arXiv:1010.3382v1 [astro-ph.GA] 16 Oct 2010

On the dynamics of tidal streams in the Milky Way galaxy

Andrew M Eyre
Rudolf Peierls Centre for Theoretical Physics

Merton College
University of Oxford

A thesis submitted for the degree of Doctor of Philosophy
at the University of Oxford

Hilary Term 2010

On the dynamics of tidal streams in the Milky Way galaxy

Andrew M Eyre
Rudolf Peierls Centre for Theoretical Physics

Merton College
University of Oxford

Abstract

We present a brief history of Galactic astrophysics, and explain the origin of halo substructure in the Milky Way Galaxy. We motivate our study of the dynamics of tidal streams in our Galaxy by highlighting the tight constraints that analysis of the trajectories of tidal streams can place on the form of the Galactic potential.

We address the reconstruction of orbits from observations of tidal streams. We upgrade the geometro-dynamical scheme reported by Binney (2008) and Jin & Lynden-Bell (2007), which reconstructs orbits from streams using radial-velocity measurements, to allow it to work with erroneous input data. The upgraded algorithm can correct for both statistical error on observations, and systematic error due to streams not delineating individual orbits, and given high-quality but realistic input data, it can diagnose the potential with considerable accuracy.

We complement the work of Binney (2008) by deriving a new algorithm, which reconstructs orbits from streams using proper-motion data rather than radial-velocity data. We demonstrate that the new algorithm has a similar potency for diagnosing the Galactic potential.

We explore the concept of Galactic parallax, which arises in connection with our proper-motion study. Galactic parallax allows trigonometric distance calculation to stars at 40 times the range of conventional parallax, although its applicability is limited to only those stars in tidal streams.

We examine from first principles the mechanics of tidal stream formation and propagation. We find that the mechanics of tidal streams has a natural expression in terms of action-angle variables. We find that tidal streams in realistic galaxy potentials will generally not delineate orbits precisely, and that attempting to constrain the Galactic potential by assuming that they do can lead to large systematic error. We show that we can accurately predict the real-space trajectories of streams, even when they differ significantly from orbits.

A thesis submitted for the degree of Doctor of Philosophy
at the University of Oxford

Hilary Term 2010

Foreword

The journey of this thesis has been a long one. Although I present thanks overleaf to those who have been of direct assistance in my work, I would like to use this page to speak of those who made the journey pleasurable as well as possible.

Firstly, I will mention my guide along the way, James Binney, without whom this endeavour would never have been started, let alone finished. While I was contemplating where to go for my PhD, a very esteemed astrophysicist, upon learning of my list of prospective supervisors, said: "James Binney is virtually in a class of his own ... in terms of the breadth and profundity of his physics knowledge. If you have the chance to work with him, I would advise you to take it." It does not take a prolonged conversation with James to confirm this opinion as abundantly true. It has indeed been a rare honour and a privilege to work with him.

Without my fellow travellers from Room 2.9 of the Peierls Centre, office life would have been so much duller. They are, in order of appearance: Ralf Donner, whose Teutonic expressivity is surely unmatched by anyone; Sarah Sägesser, who was always in the office, no matter how early I arrived; Rachel Koncewicz, whose delicious coffee gave life to the office, both metaphorically and literally; Mimi Zhang, whose inscrutable Confucianism made our chats so challenging and yet so interesting; Mike Williams, whose track record of involvement in creative projects—including but not limited to: a movie, a record label, and at least three PhD projects—is utterly bewildering; Ben Burnett, whose repartee is eclipsed only by his skill in physics; Calum Brown, who always had to hand a bolus of healthy Scottish realism; Francesco Fermani, whose fiery Italian spirit was a welcome new broom; and Alfred Mallet, who finally relieved me of the burden of being the only turncoat around.

From elsewhere in the Physics Department, I wish to mention: John Magorrian, whose ever-welcome humour is lost on many of the Peierls Centre residents, I am sure; Carole Jordan, who can be relied upon to find the silver lining in any cloud; Paul McMillan and Christophe Pichon, who both helped enspiritualize our Monday lunchtime meetings; Will Newton, whose unerring knack for attracting the unusual led to a high-speed bus chase along the M40, a well-timed escape from an Eton man waving a fistful of £20 notes, and numerous other close shaves; Garrett Cotter, who has an inexhaustible supply of tales of tutorial-room woe; and Tom Mauch, for the much-vaunted musical education that is hopefully soon to materialize!

External to the department, I have counted amongst my friends: Maja Starcevic and Kreso Petrinec, Claire Labrousse, Daniel Rotenberg and Merav Haklai-Rotenberg, Jonathan Riley, Jim Naughton, and countless others that I apologize for unfairly forgetting to mention. Whilst at home, my enduring memory is of the abundant worldly wisdom of Charlie being punctuated by the incessant banter of Clo. I would also like to thank my family, who have always given me their unconditional love and support.

Lastly, I wish to speak of Michal. She is wise and beautiful. If Michal has taken of me but a fraction of that which I have taken of her, she must indeed be the richest woman in the world. For without a shadow of a doubt in my mind, there are none that are richer than me for having shared in her life for these past few years.

Andy Eyre
Oxford, July 2010

To my mother and father

and to Michal

Acknowledgements

I would like to extend my thanks to the following, each of whose aid has directly contributed in some way to this work: Prof. James Binney, for his knowledge, patience, advice and support, and for copious quantities of his time; the members of the Oxford Dynamics group, for comradeship and helpful critique; the anonymous referees of my papers, for their helpful remarks, many of which advanced the work of this thesis; Prof. Andy Gould, whose remarks on the uncertainties in Galactic parallax calculations prompted the analysis of §4.3; Sergey Koposov for his provision of the data for Fig. 4.7; Michal Molcho for her eagle-eyed proofreading of this manuscript; and the examiners of this thesis, Vasily Belokurov and John Magorrian, for their comments and suggestions.

The following people kindly provided code that aided my calculations. I would like to acknowledge their contributions and thank them: Prof. James Binney, whose contributions included code for: the calculation of action-angle variables in the isochrone potential, the calculation of actions in Stäckel potentials, the sampling of King model distribution functions, the reconstruction of orbital trajectories using radial velocity data, and many other miscellaneous routines; Carlo Nipoti, who provided the FVFPS tree code, and updates thereof; and Paul McMillan, who provided a shake-and-bake version of Prof. Walter Dehnen's FALPOT potential calculation routines.

I acknowledge the receipt of a PPARC/STFC award during the preparation of this work.

Derivative publications

The following publications have arisen out of the work of this thesis.

Parts of Chapter 2 appeared in:
Eyre, A., & Binney, J. 2009, Mon. Not. R. Astron. Soc., 400, 528

Parts of Chapter 3 appeared in:
Eyre, A., & Binney, J. 2009, Mon. Not. R. Astron. Soc., 399, L160

Parts of Chapter 4 appeared in:
Eyre, A. 2010, Mon. Not. R. Astron. Soc., 403, 1999

Contents

1 Introduction **1**

1.1 A brief history of galactic astronomy . . . 2
1.2 Galactic cannibalism in action: tidal streams . . . 5
1.3 The use of tidal streams as probes of the potential . . . 8
1.4 Overview of this thesis . . . 10

2 Finding the orbits delineated by tidal streams **13**

2.1 Introduction . . . 13
2.2 Reconstructing orbits from tracks on the sky . . . 14
2.3 Specifying stream tracks . . . 21
2.4 Identifying dynamical orbits . . . 21
2.5 Testing the method . . . 28
2.6 The effect of changing the potential . . . 35
2.7 Conclusions . . . 38

3 Fitting orbits to streams using proper motions **43**

3.1 Introduction . . . 43
3.2 Reconstructing orbits with proper motions . . . 44
3.3 Tests . . . 49
3.4 Conclusions . . . 51

4 Galactic parallax for the tidal stream GD-1 **53**

4.1 Introduction . . . 53
4.2 Galactic parallax . . . 56
4.3 Uncertainty in Galactic parallax calculations . . . 56
4.4 Practicality of Galactic parallax as a distance measuring tool . . . 60
4.5 Tests of Galactic parallax . . . 63
4.6 Distance to the GD-1 stream . . . 65
4.7 Conclusions . . . 68

5 The mechanics of streams **72**

5.1 Introduction . . . 72
5.2 The formation of streams from tidally stripped clusters . . . 75
5.3 Stream formation in spherical potentials . . . 78
5.4 Mapping streams from action-angle space to real space . . . 89
5.5 The consequences of fitting orbits . . . 100

5.6 The action-space distribution of disrupted clusters . . . 104
5.7 Non-spherical systems . . . 116
5.8 Conclusions . . . 140

6 Conclusions **149**
6.1 Overview . . . 149
6.2 This work in context . . . 151
6.3 Future work . . . 153

A Some results in Stäckel potentials **159**
A.1 Computing the actions . . . 159
A.2 Computing the frequencies . . . 160
A.3 Computing the derivatives of the frequencies . . . 161
A.4 Computing the angles . . . 164

List of Tables

2.1 Highlighted orbits from Chapter 2 18
2.2 Configurations of test pseudo-data sets 28

5.1 Parameters for the spherical potentials used in Chapter 5 81
5.2 Highlighted orbits in the potentials of Table 5.1 82
5.3 Coordinate extrema of selected orbits from Fig. 5.5 87
5.4 Details of the cluster models used in §5.6.1 106
5.5 Parameters for the Stäckel potentials used in Chapter 5 120
5.6 Coordinate extrema of selected orbits from Fig. 5.26 123
5.7 Coordinate extrema of selected orbits from Fig. 5.27 125
5.8 Highlighted orbits in the potentials of Table 5.5 125

List of Figures

1.1 The "field of streams" of Belokurov et al. (2006) 7

2.1 Orbital energy conservation for reconstructions using perfect and erroneous data 17
2.2 On-sky projection and line-of-sight velocities for an N-body simulation of the Orphan stream, from two perspectives 19
2.3 Baseline input for the on-sky stream track of pseudo-data sets PD2–PD7 19
2.4 Baseline input for the line-of-sight velocities of pseudo-data sets PD2–PD4 20
2.5 Values of the diagnostic D' for reconstructions using PD1 28
2.6 Baseline input for the line-of-sight velocities of pseudo-data sets PD5–PD6 29
2.7 On-sky projection and line-of-sight velocities of selected orbits reconstructed from data sets PD2 and PD6 30
2.8 Heliocentric distance and tangential velocity of the best reconstruction from Fig. 2.5 . . . 31
2.9 As Fig. 2.5 but for input data sets PD2–PD7 32
2.10 Heliocentric distance for selected reconstructed orbits from the data sets PD2–PD7 33
2.11 Tangential velocities for selected reconstructed orbits from the data sets PD2–PD7 34
2.12 As Fig. 2.5 but for data set PD2 analyzed in a Kepler potential and a rising rotation-curve potential 36
2.13 Behaviour of the diagnostic D' with variation in halo mass parameter 37

3.1 Diagrams of velocities in the grf and the hrf 45
3.2 Test input for the algorithm 49
3.3 Orbital energy conservation for reconstructions using proper motions 50

4.1 Diagram illustrating the Galactic parallax effect 55
4.2 Galactic parallax distances computed using erroneous pseudo-data 62
4.3 The K10 orbit for the GD-1 stream, and pseudo-data derived from it 64
4.4 Galactic parallax distances computed from the pseudo-data of Fig. 4.3 64
4.5 On-sky position data for the GD-1 stream 65
4.6 Galactic parallax distances for the GD-1 stream 66
4.7 Galactic rest-frame proper motions for the GD-1 stream 67

5.1 In-plane rotation curves for selected potentials from Chapter 5 82
5.2 Real-space trajectories for selected orbits in the spherical potentials of Table 5.1 83
5.3 A stream formed from the orbit K1 84
5.4 Angle-space configuration for the particles shown at apocentre in Fig. 5.3 84
5.5 Details of stream geometry in the isochrone potential of Table 5.1 87

5.6 Angle-space distribution of a stream formed from the orbit I1 . . . 88
5.7 Real-space configuration for the stream shown in Fig. 5.6 . . . 89
5.8 Mapping of action-space to angle-space for the orbit I1 . . . 91
5.9 Three trajectories in angle-space, derived from the orbit I1 . . . 92
5.10 Real-space trajectories of the lines shown in Fig. 5.9, phase-matched near apocentre . . . 93
5.11 As Fig. 5.10, but phase-matched near pericentre . . . 93
5.12 As Fig. 5.10, but phase-matched at a point well away from apsis . . . 94
5.13 A stream of test particles to check that the mapping from angle-space to real-space is accurate . . . 95
5.14 Demonstration of the effect of changes in action J_r on the real-space orbital trajectory of streams . . . 98
5.15 Demonstration of the effect of changes in angular momentum L on real-space orbital trajectory . . . 99
5.16 Plots of goodness-of-fit for stream pseudo-data to the best-fitting orbit, for a range of potential parameters . . . 102
5.17 Plots of goodness-of-fit for stream pseudo-data to the best-fitting orbit, versus potential mass parameter . . . 103
5.18 Plan views of the orbits used in §5.6 . . . 105
5.19 Action-space distribution of the cluster model C1, at different times along the orbit I4 . . . 107
5.20 Radial velocity versus tangential velocity for the C1 cluster model, on the I4 orbit . . . 108
5.21 Diagram showing the Lagrange points in the reduced three-body problem . . . 110
5.22 Particle-density plots for selected panels from Figs. 5.19 and 5.23 . . . 111
5.23 Action-space distribution of particles for different cluster models on various orbits . . . 113
5.24 Angle-space and real-space distribution of particles for cluster C1 on orbit I4 . . . 115
5.25 Contours of equal density and potential, for the Stäckel models used in Chapter 5 . . . 121
5.26 Details of stream geometry for the Stäckel potential SP1 . . . 122
5.27 Similar to Fig. 5.26, but for the Stäckel potential SP2 . . . 124
5.28 Real-space trajectories for selected orbits in the Stäckel potentials of Table 5.5 . . . 126
5.29 Action-space distribution for the simulated cluster C5, on the orbit SO1 in the Stäckel potential SP1, at various points in time . . . 128
5.30 Particle-density plots for selected panels from Fig. 5.29 . . . 129
5.31 Angles for the N-body cluster shown in the bottom panels of Fig. 5.29 . . . 131
5.32 Real-space configuration of the cluster model C6 evolved on the orbit SO1 . . . 132
5.33 Actions for the cluster C6 on orbit GD1 in the Stäckel potential SP2 . . . 133
5.34 Particle-density plots for the scatter plots of Fig. 5.33 . . . 134
5.35 Angle-space configuration of the simulated GD-1 stream, shown in Fig. 5.33 . . . 135
5.36 Real-space configuration of the simulated GD-1 stream, shown in Fig. 5.33 . . . 136
5.37 Actions for the N-body cluster C7 on orbit OS1 in the Stäckel potential SP2 . . . 137
5.38 Particle-density plots for the scatter plots of Fig. 5.37 . . . 138
5.39 Angle-space configuration of the simulated Orphan stream, corresponding to the actions shown in Fig. 5.37 . . . 138
5.40 Real-space configuration of the simulated Orphan stream . . . 139

Chapter 1

Introduction

The study of astronomy has long held an ennobled position amongst the fields of natural enquiry, of which it is undoubtedly one of the oldest: the first written records of astronomical measurement were made by the ancients in the city-state of Babylon.[1] The study almost certainly goes back further, however, since many of the Babylonian constellations were named in a Sumerian dialect, and the Sumerian civilization had already crumbled by the 3rd millennium BC. The origins of Sumer predate that by some thousands of years, and are lost is the mists of prehistory: one can perhaps imagine our Mesopotamian ancestors staring at the heavens at night, and being amazed by both the regularity and spectacular beauty of the slow procession of the bodies held therein.

Even within the irrational world-views held by the ancients, it was recognized that the processes governing the movement of the heavenly bodies must be mechanical. This insight lead to the earliest recorded attempts to impart mechanical descriptions on natural phenomena. For instance, the regularity of the diurnal motion of the planets and stars about the Earth's axis led to the quantization of the passage of time: it is no coincidence that, even today, our measure of time is fundamentally a measure of angle, and indeed was without alternate physical basis until as recently as 1967.[2]

It is therefore unsurprising that history credits astronomy with provoking a most extraordinary series of physical discoveries. The motions of the planets, as observed by Tycho Brahe and resolved into orbits by Johannes Kepler, both inspired Isaac Newton and provided him with the necessary data to inform his deduction of the laws of classical mechanics and of universal gravitation. In the process of formulating his theories of planetary motion, Newton discovered the differential calculus—a necessary tool for his

[1] Perhaps competing for the title of oldest is the field of medicine, for which the *Edwin Smith Papyrus*, dated to the 16th century BC, contains rational descriptions of injury and prognosis. By comparison, the earliest records of astronomical observation are the Babylonian tablets *Enuma Anu Enlil*, also dating from sometime in the 2nd millennium BC, but which are unfortunately far from rational, since they clearly claim to have been made for the purposes of conducting magic. The earliest known rational attempt to explain astronomical phenomena is attributed to Plato's student, Eudoxus of Cnidus, who lived in the 4th century BC.

[2] The Thirteenth General Conference on Weights and Measures, 1967.

task—and consequently founded physics as a mathematical discipline. Newton then went on to write his second great treatise *Opticks*, which made great advances in the geometric description of light and directly derived from experiments he began in order to construct a better telescope.

Newton's theories were a well-spring of progress in physics, and much of this progress came in pursuit of astronomy. Although fully-formulated in *Principia* by Newton, classical mechanics later evolved under the genius of the likes of Joseph-Louis Lagrange, Pierre-Simon Laplace and William Hamilton. Lagrange and Hamilton each reformulated Newtonian mechanics into their respective, eponymous dynamics. These new dynamics readily admit solutions to problems that are computationally taxing using Newtonian mechanics, and both were directly inspired by their namesake's desire to solve problems in celestial mechanics. As too was Laplace. Intrigued by celestial mechanics, he developed potential theory, culminating in his eponymous equation: he then invented spherical harmonics in order to solve it. Laplace also laid the modern foundations of probability theory and statistics, in order to better interpret incomplete astronomical observations.

The list could go on. Astronomy has indeed inspired many of the greatest scientists in their greatest work. It is yet more remarkable, therefore, that it was only with the middle-half of the 20th century that confirmation was finally made that there existed galaxies other than our own.

1.1 A brief history of galactic astronomy

The idea of the Milky Way—easily visible on a dark night away from the glare of artificial city lighting—as a collection of stars similar to our own Sun is very old, dating from antiquity.[3] The confirmation of this supposition came by way of the persecuted genius Galileo Galilei whose self-made optical telescope allowed him to observe that the Milky Way, which appears nebulous to the naked eye, is actually comprised of countless distinct and individual stars.

The first mention of objects that we now know to be galaxies external to our own came somewhat later than did those of the Milky Way. The Magellanic clouds were first recorded by 10th century Arabian astronomers, and the knowledge of their existence was eventually brought to Europe following the global circumnavigation of the 16th century explorer Ferdinand Magellan.

It would require Newton's law of universal gravitation in order to begin to understand the true nature of galaxies, although Newton himself apparently failed to make any significant progress in the task.[4] The

[3]The Greek philosopher Democritus, a contemporary of Aristotle and Plato, was the first to be recorded espousing this view, which was perhaps informed by his strong atomist tendencies. Little of Democritus' own work survives, but we know of his views on galactic astrophysics by means of Plutarch, in *De placitis philosophorum*.

[4]Newton was traumatized by his inability to show that the Solar system was stable, since his limited calculations of the Sun-Jupiter-Earth three-body problem led him to predict the rapid ejection of the Earth from the system. Newton repaired this problem by hypothesizing the intervention of God to reset any Jovian anomalies induced in Earth's orbit. Newton's cosmology was one of an infinite constellation of static, but mutually gravitating stars: this configuration is actually unstable, and this fact was known to Newton, but he had no qualms in invoking divine intervention to maintain

English astronomer Thomas Wright (1750) was the first to understand that the Milky Way might be a rotationally-supported, flattened disk of stars, which we view as a tract across the sky by nature of our position within it. His speculation was made with little evidence to back it up. Wright also speculated that the mysterious "nebulae," such as the Magellanic clouds, might well be separate galaxies in their own right. His ideas were taken up and promulgated by Immanuel Kant, who called the structures "island universes."

Remarkably, it was not until the 1920s that observational technology improved to the point where either of these hypotheses could be conclusively confirmed. In 1924, while working at the Mount Wilson Observatory, the American astronomer Edwin Hubble used the brand-new Hooker Telescope to resolve individual stars in several nearby galaxies, including M31. Several of the stars he observed were Cepheid variables, for which the Harvard astronomer Henrietta Leavitt had some years earlier deduced a tight period-luminosity relation. With their intrinsic brightness determined by this relation, the faint apparent magnitude of the stars put them far beyond the most generous estimates for the extent of the Milky Way galaxy at that time (Hubble, 1926). M31 and its companion nebulae could be nothing other than external galaxies. Thus, the Wrightian conjecture of "island universes" made almost 200 years earlier was proved and the study of extragalactic astronomy was born.

As has often been the case in the history of our field, fate conspired that the mathematical tools and the experimental machinery to probe a new area of science became available at the same time. Einstein's general theory of relativity (1915) had provided the physical framework upon which consistent cosmological theories could be constructed. Hubble was again instrumental in advancing the field, and in 1929 he made the discovery of a linear relationship between the distance to, and the measured line-of-sight velocity of, far-away galaxies (Hubble, 1929). This was the first observational evidence for the expansion of the Universe, and the modern study of cosmology was born, with Hubble having launched his second new field of natural enquiry in about as many years.[5]

The consequences of Hubble's two great discoveries were enormous. The expansion of the Universe

it indefinitely, just as he had for the Solar system. Newton's unwarranted belief in a static universe puts him good company. Albert Einstein's similarly-unwarranted belief in the same led him to insert an otherwise unprovoked constant of integration—the cosmological constant—into his general-relativistic field equations. Einstein later regretted this action, calling it "the biggest blunder of (my) life" (Gamow, 1970). The modern understanding of the Newtonian Solar system shows that it is indeed subject to rapid disintegration on account of Jovian perturbations: it is of some irony that Einstein's own general-relativistic corrections to Mercury's orbit are required in order to detune the resonance between Jupiter's orbit and Mercury's and thus stabilize the system, which would otherwise result in the ejection of Mercury, followed by the other inner planets in a few million years (Laughlin, 2009). If only Newton had known general relativity!

[5]Infamously, Hubble never won the Nobel Prize in Physics for his groundbreaking contributions to our understanding of the Universe. At the time, astronomy was not considered in the remit of the physics Prize, and despite a growing clamour in the scientific community for the Prize to be awarded to Hubble, he died suddenly in 1953 having not received it. Astronomy and astrophysics were admitted to the list of eligible fields of study that very year. Indeed, at the time of his death—but obviously unknown to Hubble himself—the 1953 Nobel Prize committee, amongst whose august members were counted Enrico Fermi and Subrahmanyan Chandrasekhar, had already unanimously voted Hubble to receive that year's physics Prize (Soares, 2001). Nobel Prizes cannot be awarded posthumously: the first Nobel Prize awarded for an astrophysical discovery eventually went to Hans Bethe in 1967, for his explanation of the nuclear fusion processes that power stars (Bethe, 1972).

requires that at some time in the past all matter was coincident, and hence the Universe is of finite age. Although it would take many decades before this age was known with any kind of precision, it was apparent from Hubble's first observations that the age of the entire Universe was not incomparable to the geological age of the Earth—some billions of years.

One conclusion to be drawn from this is that galaxies are not and were never steady-state objects. Since they cannot be substantially older than the stars within them, and given the great distance scales that they span, galaxies cannot be dynamically very old. Indeed, our own Milky Way cannot have completed more than 70 complete revolutions since the big bang, assuming that it formed shortly thereafter.

It thus became—and remains—a core problem in modern astrophysics to explain the structure, formation and evolution of galaxies. It was immediately possible for Hubble's contemporaries, such as James Jeans and Arthur Eddington, to begin this difficult task because the evocation of statistical mechanics by James Clerk Maxwell and Ludwig Boltzmann in the late 19th century had provided the tools necessary to begin to understand the bulk motion of stars under the influence of gravity. In combination with orbital mechanics, this laid the framework for the study of stellar dynamics, which deals with systems comprised of many-fold more bodies than does celestial mechanics, its direct intellectual predecessor.

Here we will leave our whistle-stop history tour of galactic astronomy, but for one small diversion of direct relevance to our work. Of the many astounding scientific discoveries of 20th century, one of the least expected, and certainly one of the least well-understood, has been the growing—and by now, colossal—body of evidence that there is simply insufficient baryonic matter in most galaxies to explain the observed kinematics.

The first observation to this effect was reported by the Swiss astronomer Fritz Zwicky (1933), who combined the virial theorem with measurements of velocity dispersion in galaxy clusters to argue that the bulk of their mass must be in non-luminous matter. Further progress awaited technological advances. Up until the late 1950s, galactic astronomy consisted mostly of observations of stars at optical or near-optical wavelengths. The discovery of the so-called 21-cm neutral hydrogen transition, which arises because of spin-spin coupling between electron and proton in the ground state of atomic hydrogen, heralded yet another revolution in galactic astronomy. Radio astronomy observations of 21-cm emission allowed the optically-transparent, non-ionized gaseous content of galaxies to be mapped for the first time, and since then with ever increasing sensitivity.

The first 21-cm maps of gas in external galaxies showed that the circular speed of matter was independent of radius. One corollary of this surprising observation is that the mass interior to a given radius must grow linearly with radius: thus, the dynamical mass in these galaxies was not concentrated in those regions of high luminous mass, as had previously been supposed. Even more significantly, the 21-cm maps were used to quantify the mass of gas in these galaxies, which was then added to star

counts to estimate the total content of baryonic matter. The resulting estimate for the mass content of these galaxies was insufficient to explain the observed kinematics of the gas and stars, with this so-called "missing mass" problem becoming most acute at large radii where the density of baryonic matter falls rapidly, but where the kinematics indicate that most of the matter is concentrated (Binney & Merrifield, 1998, §8.2.4).

The description of this "dark matter" is one of the foremost problems in modern astrophysics, and neither the formation nor the evolution of galaxies can be properly studied without addressing it. N-body simulation of cosmological structure formation has provided clues as to what the dark matter distribution should look like (Navarro et al., 1997), and the suggested profiles have some corollary in observations of external galaxies (e.g. Rix et al., 1997). However, the universal applicability of the simulation results is far from proved, and apart from its general presence, the distribution of dark matter in the Milky Way in particular is still not well-determined (Smith et al., 2007).

Observations are needed. Unfortunately, all attempts to directly detect particle annihilation signatures from concentrations of dark matter have as yet been unsuccessful (Ahmed, 2009), and in any case, the detection rate of such signatures is unlikely to ever be high enough for useful imaging to take place. Our only option to examine the dark matter content of the Milky Way is to utilize that very mechanism by which we hypothesize its existence in the first place, namely, the effects of its gravity on the dynamics of luminous, observable matter.

It will therefore be the topic of this thesis to address certain indirect methods of probing the mass distribution of our Galaxy. To do this, we will examine the mechanics of tidal streams.

1.2 Galactic cannibalism in action: tidal streams

"Galactic cannibalism"—to use the words of Ostriker & Hausman (1977)—refers to the merging of individual galaxies to form a greater whole and has been increasingly recognized over the last 30 years as playing a significant role in determining the structure of galaxies in general, and the Milky Way in particular (Binney & Tremaine, 2008, §8). Indeed, it is a fundamental tenet of the White & Rees (1978) model of galaxy formation, which has dominated our understanding since its inception. In this picture, baryonic components are embedded in massive dark haloes, which cluster and then merge purely under the force of gravity. Radiative processes then cause the gaseous phase to cool and contract and eventually form stars at the bottom of the potential well.

In support of this model, we note that collisions between external galaxies have been known about for a long time. The M51 "Whirlpool" and NGC 4038/9 "Antennae" galaxies are all undergoing obvious merging events, and all were discovered in the 18th century, although their nature was not understood

at the time. Violent mergers, such as those of the Antennae, are often associated with tidal tails: high-energy ejecta made up of stars and gas, catapulted from the edges of the colliding objects by immense tidal forces, and only marginally bound, if bound at all, to the resulting combined host mass (Toomre & Toomre, 1972). Many more such merging systems are now known, and these interactions are believed to be commonplace.[6]

The hierarchical galaxy formation model is particularly successful in explaining the presence of the Milky Way's halo: a spheroid of old, metal-poor stars and globular clusters, extending some tens of kpc from the Galactic centre (Binney & Merrifield, 1998, §10.5). The halo contains very little gas and dust and is hard to see how the halo stars could form in situ. The White & Rees (1978) mechanism answers this, by explaining the halo as the phase-mixed stellar remnant of long-ago mergers.

However, up until the 1990s, very little evidence of cannibalization of its satellites by the Milky Way had been seen at all: the Large Magellanic Cloud was identified by Mathewson et al. (1974) as losing mass to the Milky Way halo, but the mass lost is in gas and not stars, and rather than being evidence of a gravity-driven merging event, it is most likely that the observed 'tail' is simply the streamlined wake of the Cloud's gaseous envelope being stripped due to ram pressure from the Milky Way's own halo gas (Moore & Davis, 1994).

The advent of multi-million particle N-body simulations of cosmological structure formation allowed quantitative predictions for the expected number of Milky Way satellite galaxies and merger remnants to be made (Navarro et al., 1997; Moore et al., 1998; Bullock et al., 2001). The outcome was problematic: simulations showed that, although the numbers of high-mass satellites was predicted almost perfectly, the Milky Way ought to have accreted an order of magnitude more low-mass satellite galaxies than were actually known at the time (Klypin et al., 1999). Attempts were made to repair the "missing satellite" problem by proposing mechanisms to shut off star formation in low-mass haloes, thus rendering them invisible, but the situation still remained highly unsatisfactory (Klypin et al., 1999; Bullock et al., 2000; Moore et al., 2006).

The predictions of the simulations went further. Long-lived stellar substructure in the Milky Way's halo was shown to be a consequence on ongoing merger activity, and this substructure was associated with kinematic and chemical signatures that ought to be observable (Bullock & Johnston, 2005). Indeed, the substructure was reckoned to be permeate the Galaxy with sufficient density to leave kinematic traces in the Solar neighbourhood (Helmi & White, 1999). In this way, some of the first direct evidence for substructure resulting from past mergers was identified by Helmi et al. (1999) using data from the

[6]Such a merger is forecast between the Milky Way and M31, to take place in about 3 billion years time (Cox & Loeb, 2008). The structure of both galaxies will be destroyed, and a massive elliptical galaxy will emerge in their stead, although it is most unlikely that individual solar systems will be directly affected by the merger. The elderly Sun will still be in the main sequence at this time: for whoever or whatever life inhabits the Earth, the spectacle in the night sky will be extraordinary.

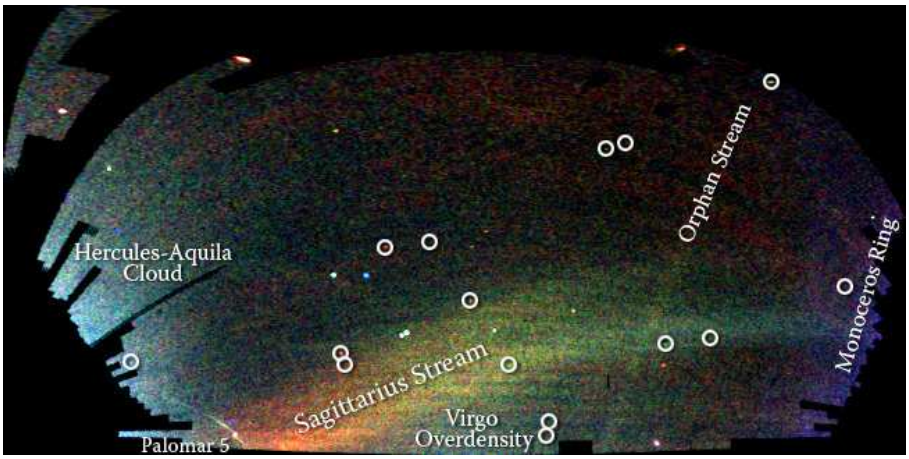

Figure 1.1: The "field of streams" of Belokurov et al. (2006). The observation of tens of millions of stars with SDSS, combined with appropriate cuts to the data, expose the dramatic merger-history substructure present in the Milky Way's halo. The colours are for r magnitude, which proxies for heliocentric distance in this image. The circles highlight clusters/galaxies newly discovered in this image. *Credit: V. Belokurov and the Sloan Digital Sky Survey. Image source: sdss.org.*

Hipparcos satellite (van Leeuwen, 2007).

The discovery of the Sagittarius Dwarf galaxy by Ibata et al. (1995), on the far side of the Milky Way, provided dramatic evidence for ongoing merger activity. This galaxy stands out amongst the other known companion galaxies on account of both its high mass and close proximity to the centre of the Milky Way. Dynamics requires the Milky Way to impart strong tidal forces across the Sagittarius Dwarf galaxy, with it being so large and so close, and thus the existence of a tidal stream of stripped dark matter and stars was predicted soon after its discovery (Velazquez & White, 1995; Johnston et al., 1995). It took nearly 10 years before the stars of the massive Sagittarius stream were conclusively observed by Majewski et al. (2003) in infra-red data from the Two Micron All-Sky Survey (2MASS, Skrutskie et al., 2006).

It is difficult to observe merger substructure in external galaxies, because it phase-mixes rapidly, and is only detectable thereafter in the form of kinematic and chemical signatures, which require observations of such precision that they can only be performed in the Milky Way. It was therefore with some great anticipation that data from the automated Sloan Digital Sky Survey (SDSS, York et al., 2000) arrived. This project used a purpose-built 2.5-m telescope to survey tens of millions of stars, as faint as magnitude $r \sim 24$, away from the Galactic plane. Nonetheless, despite its impressive performance, the very faint substructure of the Galactic halo still required careful data processing in order to expose its signal above the noise of the halo stars.

Fig. 1.1 shows one result of the effort: an on-sky map along the north Galactic pole, which was reported by Belokurov et al. (2006). This image shows the halo to be criss-crossed by extant tidal streams, and dotted with hitherto undetected low-mass halo-dwelling galaxies. Indeed, the observational

evidence unearthed by the SDSS, and its follow-on extension programme (SEGUE, Yanny et al., 2009), revealed a dramatic number of tidal streams in the Milky Way halo, often with no obvious associated progenitor object (Odenkirchen et al., 2003; Majewski et al., 2003; Yanny et al., 2003; Belokurov et al., 2006, 2007; Grillmair, 2006a; Grillmair & Dionatos, 2006; Grillmair & Johnson, 2006; Grillmair, 2009; Newberg et al., 2009).

These streams are almost certainly the remnants of objects of extragalactic origin. Short of a major merger, it is hard to envisage a scattering event that could launch an existing Milky Way globular cluster onto a lower-energy orbit on which it cannot survive. Hence, it is likely that these streams originated either from former members of the cluster system of a larger cannibalized galaxy, which could not then survive on their new orbit around the Milky Way, or tiny galaxies whose streams are the fading echo of their cannibalization by the Milky Way in their own right.

The nature of the progenitors of these streams may be important for tracing the merger history of our Galaxy (Helmi, 2008). However, in this thesis, we will be concerned with the use of tidal streams as probes of the matter distribution of the Milky Way itself. As such, the details of their origin do not concern us greatly: we simply accept that at some point in their history, these objects have found their way onto orbits around our Galaxy, along which they experience tidal forces strong enough to promote their disintegration.

1.3 The use of tidal streams as probes of the potential

The possibility that the morphology of tidal tails from colliding galaxies could act as a probe of the potential has been recognized for almost as long as the physics of the tails themselves has been understood (Faber & Gallagher, 1979), although it is only much more recently (Dubinski et al., 1996) that simulation technology has been sufficient to make a serious attempt at constraining the potential with models of galaxies undergoing a major merger.

However, on account of the great disparity in mass between the Milky Way and its victims, the mergers observed are less violent than those seen between objects of similar mass. We should therefore be careful to distinguish the tidal *streams* resulting from the simple disintegration of clusters and small galaxies in the gravity field of a massive host—as is typically seen in the Milky Way—from those tidal *tails*, which result from major mergers. The stars of the former remain in low-energy orbits around the massive host, while the stars in the latter are unbound, or nearly so.

Furthermore, although the same fundamental physical principles underlie the creation of both species, the tidal streams that result from cannibalization of low-mass satellites, such as we see in the Milky Way halo, are almost certainly more useful than tails from major mergers in probing the potential of an

individual galaxy. The reason is that major mergers must involve the catastrophic agglomeration of the dark matter haloes, wiping out all precursory substructure in favour of the newly virialized dark halo blob. Indeed, in the White & Rees (1978) model of galaxy formation it is the very occurrence of this dark-matter agglomeration which seals the fate of the merger. Hence, any probe of the dark matter distribution resulting from such a merger would be of comparatively little use in constraining the dark matter distribution of the precursor objects. Lastly, the unbound or marginally-bound tail stars from a major merger feel the gravity of their former hosts only weakly, which therefore has little effect on their motion, making them less sensitive probes overall.

The use of Milky Way streams as probes of the Galaxy potential stems from a recognition of the similarity between the trajectory of a stream, and the orbital trajectory of the progenitor object (McGlynn, 1990; Johnston et al., 1996). Observing the trajectory of an orbit places strong constraints on the gravitational force-field that gave rise to that trajectory, and hence also constrains the distribution of matter that generates that field (Binney, 2008). Indeed, it is precisely the knowledge of orbital trajectories in the Solar system that allows the mass of the Sun and the planets to be so well constrained. Similarly tight constraints can be placed on the mass of the Milky Way's black hole, from the orbital trajectories of S-stars near the Galactic centre (Ghez et al., 2005).[7]

In recognition of the diagnostic power of orbits, many recent papers have put some considerable effort into the attempt to locate the orbits delineated by tidal streams (Law et al., 2005; Fellhauer et al., 2007a,b; Odenkirchen et al., 2009; Willett et al., 2009; Koposov et al., 2010). Unfortunately, success in this endeavour, to date, has been somewhat limited. The traditional methods of finding orbits consistent with the data, namely, to search over a range of initial conditions from which one integrates the equations of motion, have resulted in fewer convincing fits to the data than might be expected (Eyre & Binney, 2009a).

Part of the problem is the enormous space of initial conditions that must be searched over. Even if one is lucky enough to stray upon an initial condition which reproduces an orbit roughly consistent with the data, it is not clear that a better match could not be found, with a different initial condition, and perhaps in a different potential. Ockham's razor requires that we be able to convincingly falsify any theory of equal or lesser complexity to our own: to make a strong statement about the Galactic potential, we must be able to show that *no* other orbits in a given potential are compatible with the data.

"Geometrodynamical" methods—in the parlance of Jin & Lynden-Bell (2008)—are one answer to this difficulty. Such techniques utilize additional measurements of the stream, such as line-of-sight

[7]Both of these latter problems are somewhat better conditioned than ours, because the shape of the applicable potential—i.e. that of Kepler—is already known.

velocity measurements, to place additional constraints upon its trajectory (Jin & Lynden-Bell, 2007, 2008; Binney, 2008). These additional constraints substantially reduce the scope of the problem: solutions are now parameterized by the form of the potential and perhaps one other measurement. This substantial reduction in the space that parameterizes solutions makes an automated search over it feasible. Hence, the orbit consistent with a stream in a given potential can be isolated; or it can be shown that the data are incompatible with a given potential. This first half of the this thesis will advance the work of Jin & Lynden-Bell (2007) and Binney (2008) by specifying procedures to make such methods robust against errors in input data. We will also derive a new method, similar to the Binney (2008) procedure, but one that uses proper-motion measurement data as input instead of radial-velocity measurements.

Another issue that affects attempts to constrain the Galactic potential using streams is the degree to which the latter truly represent orbits. The belief that they do seems to originate in empiricist observations, made from the results of N-body simulations (e.g. McGlynn, 1990). Latterly, evidence has come to light, again from N-body simulations, that this belief may not be strictly true (Choi et al., 2007; Eyre & Binney, 2009b). Indeed, it turns out that the belief has no basis in classical mechanics, and quite the opposite is true: generally streams do not delineate orbits. In the second half of this thesis, we will demonstrate this from first principles. We will also examine techniques that attempt to ameliorate the use of such non-orbital stream data to constrain the Galactic potential.

1.4 Overview of this thesis

This thesis comprises several related studies in the dynamics of tidal streams around our Galaxy. The common thread that links these studies is the desire to exploit observations of tidal streams to place constraints upon the form of the Galactic potential.

The thesis is laid out in four substantive chapters according to the descriptions that follow. In addition to those, Chapter 6 reviews our findings in the context of astrophysics as a whole. We also provide some ancillary results to the calculations of Chapter 5 in Appendix A.

1.4.1 Chapter 2: Finding the orbits delineated by tidal streams

Binney (2008) and Jin & Lynden-Bell (2007) independently reported an algorithm for reconstructing full phase-space trajectories for tidal streams, given only the projection of a stream's trajectory onto the plane of the sky, and measurements for the line-of-sight velocities everywhere along the stream. In this chapter, we demonstrate that the applicability of this algorithm is limited by errors in the input tracks from the following sources: likely statistical errors from observations, and systematic errors due to the fact that streams do not precisely delineate individual orbits.

We offer a procedure to overcome these difficulties by specifying a parameter space, which describes modifications to the baseline input in a way that is likely to correct for the above-mentioned errors while still remaining consistent with specified uncertainty in the baseline input. We then describe procedures to search over this parameter space, while applying the Binney (2008) algorithm to isolate those modifications that correspond to orbits. In this way, we are able to find orbits consistent with stream observations, without being hamstrung by errors in input data.

The ultimate goal of our work is to diagnose the Galactic potential. Binney (2008) showed that precise measurements of streams could place stunningly tight constraints on the potential. We illustrate the extent to which this is possible using realistic data.

1.4.2 Chapter 3: Fitting orbits to streams using proper motions

The major limitation on the work of Chapter 2 is the lack of line-of-sight velocity measurements to distant main-sequence stream stars. Although obtaining such measurements is within the capability of the technology of the day, it does require the commitment of 8-m class telescope time, which is unlikely to be forthcoming soon for more than a few streams.

In this chapter, we present a possible alternative: the probing of the potential using proper-motion measurements along streams. Following the same logical schema as is used in Binney (2008), we develop an algorithm to reconstruct the orbits of streams using such measurements. We find that it is equally efficacious to use proper-motion measurements of streams to reconstruct orbits and constrain the Galactic potential.

1.4.3 Chapter 4: Galactic parallax

Measuring distances in our Galaxy is critical to understanding its structure. However, line-of-sight distances can typically be measured with only relatively poor precision, and this lack of precision is manifest in the most basic of Galactic parameters, for instance, the distance to the Galactic centre (McMillan & Binney, 2010).

The gold standard of Galactic distance estimation is trigonometric parallax. However, its applicability is effectively limited to nearby stars. For more distant objects, alternative techniques such as photometric distance estimation must be used. Unfortunately, such non-geometric techniques necessarily rely on assumptions about stellar chemistry and composition that add complexity and uncertainty to measurements made using them.

The work of this chapter shows that by making use of the known trajectory of stream stars on the sky, and given accurate enough proper-motion measurements, it is possible to calculate trigonometric

distances to stars in distant streams. This effect, which we call "Galactic parallax", has a range some 40 times greater than that of conventional trigonometric parallax, given similarly accurate measurements of motion on the sky. We examine in detail the practicality and the limitations of distance estimation using the effect, and we demonstrate its utility by accurately computing the distance to the tidal stream GD-1 (Grillmair & Dionatos, 2006).

1.4.4 Chapter 5: The mechanics of streams

There has been some noise recently in the literature as to whether tidal streams can be taken to precisely delineate orbits (Choi et al., 2007; Eyre & Binney, 2009a,b). Standard techniques for constraining the potential by fitting orbits to streams rely upon the assumption that they can (e.g. Newberg et al., 2010). In this chapter, we demonstrate by use of analytical mechanics that streams, in general, do not delineate orbits. We further show that constraining the potential by assuming that streams make good proxies for orbits can lead to serious systematic error.

However, we also show that with relatively simple models of the phase-space distribution of disrupted clusters, it is possible to predict perfectly the trajectory of a stream, even when this differs significantly from the trajectory of a valid orbit. Thus, we conjecture it may be possible to repair the fitting algorithms, by having them utilize such stream trajectories instead.

Chapter 2

Finding the orbits delineated by tidal streams

2.1 Introduction

In this chapter, we first recount a procedure, independently described by Jin & Lynden-Bell (2007) and Binney (2008, herein B08), that permits the reconstruction of full phase-space information for an orbital trajectory, given only the assumption of the host Galaxy potential, and knowledge of both the track of the trajectory across the sky, and the rest-frame line-of-sight velocity down the track.

In B08 it was shown that, in addition to reconstructing phase-space information, by further examining which of the reconstructed trajectories (if any) constitute dynamical orbits, it is possible to utilize such tracks to diagnose the host potential. Given good enough input data, the precision with which both the trajectory can be reconstructed, and the potential diagnosed, was shown to be exquisite: distances and potential parameters are predicted to better than one per cent, which is far better than anything that has hitherto been possible with conventional techniques.

This advent of this procedure is particularly exciting in the context of Galactic astrophysics, since tidal streams from disrupted satellite galaxies have been held for some time to effectively delineate the orbit of their progenitor (Johnston et al., 1996; Odenkirchen et al., 2003). Thus, the detailed examination of the kinematics of these structures may well prove an important method for determining the nature of the Galactic potential.

The major limitations with the procedure as detailed in B08 are two-fold. Firstly, input data derived from observations will be subject to error, which limits the degree to which they can accurately represent an orbital track. Secondly, it has been noted (e.g. Choi et al., 2007) that tidal streams do *not* delineate

individual orbits. This compromises a core requirement of the B08 procedure, since the input tracks now no longer represent dynamical orbits, and hence dynamics cannot be used to select the physical reconstructions from the unphysical ones.

The work of this chapter, much of which has been published in an article by Eyre & Binney (2009b), is devoted to examining these limitations, and exploring procedures by which they may be overcome.

The chapter is laid out as follows. §2.2 recounts some of the work of B08, upon which this chapter draws heavily, and demonstrates its limitations when applied to data for a simulated tidal stream. §2.3 briefly examines, with the aid of a simulated example, the constraints that observations of a stream's track place on the orbits of the stream's constituent stars. §2.4 describes procedures by which we identify those orbits (if any exist) that are consistent with such constraints. §2.5 tests the method, with the aid of a simulated example. §2.6 examines our ability to diagnose the host potential. In §2.7 we present our concluding remarks.

For the remainder of this chapter, the reference frame used is the inertial frame in which the Galactic centre is at rest; consequently line-of-sight velocities are obtained by subtracting the projection of the Sun's motion from the measured heliocentric velocities. We assume complete knowledge of the velocity of the Sun with respect to the Galactic centre throughout most of this chapter.

Except where stated otherwise, orbits and reconstructions are calculated using the Galactic potential of Model II from Binney & Tremaine (2008, Table 2.3), which is a slightly modified version of a halo-dominated potential described by Dehnen & Binney (1998a). We take the distance to the Galactic centre to be 8 kpc and from Reid & Brunthaler (2004) (for V and W) and Dehnen & Binney (1998b) we take the velocity of the Sun in the Galactic rest frame to be $(U, V, W) = (10.0, 241.0, 7.6)\,\mathrm{km\,s^{-1}}$.

2.2 Reconstructing orbits from tracks on the sky

The following formulation for reconstructing complete phase-space information from a single point in space, an on-sky track, and the line-of-sight velocity measurements along that track, was discovered independently by Jin & Lynden-Bell (2007) and B08. Our work draws heavily on the latter formulation, so it is that formulation which we recount here.

Consider a tidal stream around the Galaxy, which we assume to delineate an orbit. Now let $\mathbf{r}$ be the position vector from the Sun to a star in this stream, let $\mathbf{v}$ be the rest-frame velocity of that star, and let $\Phi(\mathbf{r})$ describe the potential of the Galaxy. At this point we assume the track to delineate an orbit perfectly, so its trajectory obeys the equations of motion

$$\ddot{\mathbf{r}} = -\nabla\Phi(\mathbf{r}) = \mathbf{F}(\mathbf{r}), \tag{2.1}$$

where $\mathbf{F}$ is the acceleration due to the gravity of the Galaxy. We define v_r as the radial component of velocity,

$$v_r = \hat{\mathbf{r}} \cdot \mathbf{v} = \dot{r}, \tag{2.2}$$

and we note that its derivative

$$\dot{v_r} = F_r + \mathbf{v} \cdot \frac{\mathrm{d}\hat{\mathbf{r}}}{\mathrm{d}t}, \tag{2.3}$$

where we have defined $F_r = \mathbf{F} \cdot \hat{\mathbf{r}}$ as that component of the acceleration $\mathbf{F}$ along the line of sight. From the definition of $\mathbf{v}$ we have

$$\mathbf{v} = \frac{\mathrm{d}(r\hat{\mathbf{r}})}{\mathrm{d}t} = v_r\hat{\mathbf{r}} + r\frac{\mathrm{d}\hat{\mathbf{r}}}{\mathrm{d}t}, \tag{2.4}$$

which rearranges to

$$\frac{\mathrm{d}\hat{\mathbf{r}}}{\mathrm{d}t} = \frac{(\mathbf{v} - v_r\hat{\mathbf{r}})}{r}. \tag{2.5}$$

Combining the above expression with equation (2.3), we find

$$\dot{v_r} = F_r + \frac{\left(v^2 - v_r^2\right)}{r} = F_r + \frac{v_t^2}{r}, \tag{2.6}$$

where $\mathbf{v}_t$ is the component of the star's velocity in the plane of the sky. $\mathbf{v}_t$ must satisfy the relation

$$v_t^2 = \left(r\frac{\mathrm{d}\hat{\mathbf{r}}}{\mathrm{d}t}\right)^2 = \ (r\dot{u})^2 \equiv r^2\left(\dot{b}^2 + \dot{l}^2\cos^2 b\right), \tag{2.7}$$

where (l, b) are the on-sky Galactic coordinates and where we now fix the meaning of the parameter u to be the angular distance along the track. Combining equation (2.6) and equation (2.7), we obtain the non-linear ODE

$$\frac{\mathrm{d}v_r}{\mathrm{d}t} = F_r + r\left(\frac{\mathrm{d}u}{\mathrm{d}t}\right)^2. \tag{2.8}$$

We can rearrange this equation for $\mathrm{d}t/\mathrm{d}u$ as follows. Utilizing the chain rule and multiplying through by $(\mathrm{d}t/\mathrm{d}u)^2$,

$$\frac{\mathrm{d}t}{\mathrm{d}u}\frac{\mathrm{d}v_r}{\mathrm{d}u} = F_r\left(\frac{\mathrm{d}t}{\mathrm{d}u}\right)^2 + r. \tag{2.9}$$

This equation is quadratic in $\mathrm{d}t/\mathrm{d}u$, and it can be solved for that quantity

$$\frac{\mathrm{d}t}{\mathrm{d}u} = \frac{1}{2F_r}\left(\frac{\mathrm{d}v_r}{\mathrm{d}u} - \sqrt{\left(\frac{\mathrm{d}v_r}{\mathrm{d}u}\right)^2 - 4\,r F_r}\right), \tag{2.10}$$

where the choice of the negative root is made by requiring that $\mathrm{d}t/\mathrm{d}u$ is always positive. Equation (2.10) forms a system of coupled ODEs along with

$$\frac{\mathrm{d}r}{\mathrm{d}u} = v_r \frac{\mathrm{d}t}{\mathrm{d}u}, \tag{2.11}$$

which follows from the definition of v_r and the chain rule.

Momentarily assume that the host potential $\Phi(\mathbf{r})$ is known, and that the line-of-sight velocity $v_r(u)$ is known everywhere along an on-sky track $[l(u), b(u)]$. Given a single initial distance r_0 to some fiducial point on that track, equations (2.10) and (2.11) can be integrated. The resulting solution describes a trajectory for which full phase-space information is defined.

This result was reached independently by B08 and Jin & Lynden-Bell (2007), although the latter did not realize that only a subset of the solutions to equations (2.10) and (2.11) could be dynamical orbits. If one is to unlock the full diagnostic power of streams, it is important to isolate those solutions that are dynamical orbits. Further, if one can show that *no* solution of equations (2.10) and (2.11) with *any* initial distance r_0 is dynamical, then it follows that the assumed form for the host potential $\Phi(\mathbf{r})$ must be wrong. B08 identified as dynamical those solutions with minimal rms orbital energy variation down the track, and for a given set of input data, isolated them all by means of a comprehensive search over r_0.

2.2.1 The problem with erroneous data

B08 showed that equations (2.10) and (2.11) can locate dynamical orbits with exquisite precision, if given perfect input data in the form of $[l(u), b(u)]$ and $v_r(u)$. The example in Fig. 2 of B08 showed this for an orbit in the Miyamoto-Nagai potential (Miyamoto & Nagai, 1975). However, B08 also showed that the ability of the technique to identify dynamical orbits quickly degrades when the input data are convolved with small random errors.

Fig. 2.1 demonstrates both these effects in the more realistic Model II potential used throughout this chapter. The right panel of Fig. 2.1 shows two on-sky tracks. The black track is the projection of a segment of the PD1 Test Orbit, described in Table 2.1. The red track is derived from the black track, but the input data $[l(u), b(u)]$ and $v_r(u)$ have been modified by the addition of random fluctuations, to $b(u)$

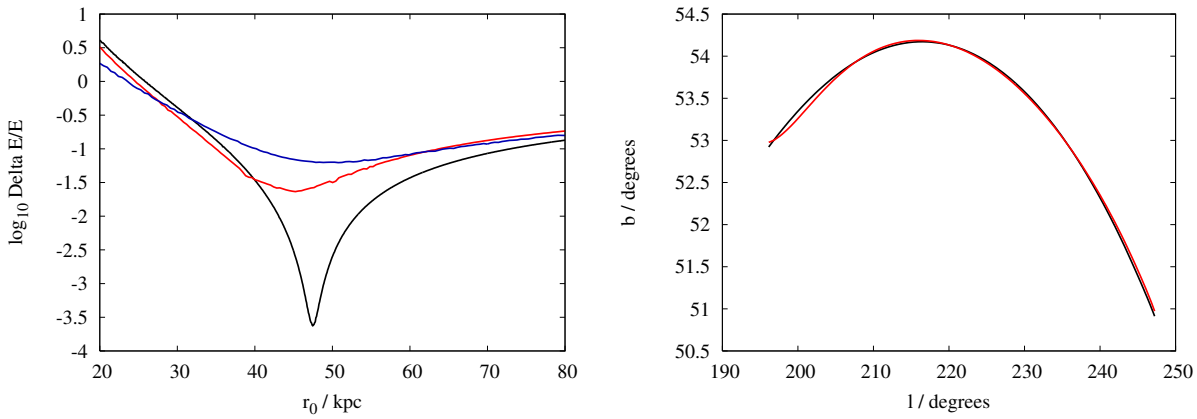

Figure 2.1: Left panel: normalized rms orbital energy, vs. initial distance to the stream r_0, for three input tracks. Right panel: on-sky plot of two tracks, used as the [l,b] input data for two of the curves in the left panel. The black curve in both panels is for the PD1 Test Orbit described in Table 2.1. The red curve in both panels is as the black curve, but with a small amount of random noise added to the input data. The blue curve is for a track derived from the simulated Orphan stream, shown in Fig. 2.2.

and $v_r(u)$, with the approximate[1] amplitude of 3 arcmin and $0.35\,\mathrm{km\,s^{-1}}$, respectively. The magnitude of these fluctuations is, respectively, smaller than the on-sky width of the narrowest known stream, and smaller than the velocity-measuring precision of an 8m-class spectrograph-equipped telescope when observing distant stream stars. Thus, such fluctuations are plausible estimates for likely errors introduced by the observational process. We note that the resulting red track is only marginally distinguishable from the black track, even at the augmented scale of the right panel of Fig. 2.1.

Using each of the tracks in the right panel, trajectories were reconstructed using equations (2.10) and (2.11) for a range of initial distances r_0. Along each such reconstructed trajectory, the normalized rms orbital energy variation

$$\frac{\Delta E}{E} \equiv \sqrt{\frac{\langle E^2 \rangle}{\langle E \rangle^2} - 1}, \tag{2.12}$$

was computed, where $E_i = v_i^2/2 + \Phi(\mathbf{r}_i)$ at a point $\mathbf{r}_i$ along the track. The numerical scheme used in the solution of equations (2.10) and (2.11) was identical to that of B08, save for a minor upgrade to the endpoints of the splines, detailed in §2.4 below.

The left panel of Fig. 2.1 shows $\Delta E/E$ versus initial distance r_0, when each of the black and red tracks is used as input. In the case of the perfect input of the black track, the correct initial distance $r_0 \simeq 47\,\mathrm{kpc}$ is identified with little scope for error, since $\Delta E/E$ reaches a deep and sharp minimum at that point. However, the effect of the random fluctuations in the input data on the ability to identify

[1]The detail of the random noise is as follows. The input data for the black track were used as the baseline input $b_b(l)$ and $v_{rb}(u)$ in the Chebyshev series of equations (2.20). The coefficients a_n and b_n, with $n = (1, 10)$, were each randomly sampled from a uniform distribution, with maximum permitted values of $0.35\,\mathrm{km\,s^{-1}}$ for the a_n and 3 arcmin for the b_n, respectively. The noisy red track is described by the $b(l)$ and $v_r(l)$ that result from equations (2.20).

Table 2.1: Parameters of highlighted orbits from this chapter. The coordinate system used is right-handed with $\hat{x}$ pointing away from the Galactic centre and $\hat{y}$ opposite the sense of Galactic rotation.

	position (x,y,z)/kpc	velocity (x,y,z)/km s^{-1}
N-body Orphan	$(28.1, -10.0, 34.0)$	$(-89.2, -37.1, -76.2)$
PD1 Test Orbit	$(35.5, 7.80, 37.8)$	$(-7.97, 56.3, 47.8)$

dynamical orbits has been very damaging. The depth of the minimum has decreased by two orders of magnitude, and the minimum has changed location to $r_0 \simeq 45\,\mathrm{kpc}$. Although perhaps the extreme distances of $r_0 \sim 20\,\mathrm{kpc}$ and $r_0 \sim 80\,\mathrm{kpc}$ could still be ruled out, there is now only a marginal basis on which to select a distance at $r_0 \simeq 45\,\mathrm{kpc}$ over another distance, since a different random fluctuation could put the easily-moved minimum elsewhere. Most disappointingly, all power to identify dynamical orbits from amongst the reconstructed trajectories has been lost, since the best trajectory only conserves energy to one part in 40, which would wash out most of the exquisite detail shown in Fig. 3 of B08, which was key to the ability of that work to diagnose the potential.

We therefore conclude that the B08 procedure does not cope well with likely observational error.

2.2.2 The problem with real streams

The previous section, and the work of B08 that it recaps, are predicated on the assumption that streams precisely delineate orbits. Recent studies involving N-body simulations of such streams (e.g. Dehnen et al., 2004; Choi et al., 2007) have made it clear that they do not. Chapter 5 of this thesis investigates the mechanics of tidal stream formation in some detail, and concludes with the ability to predict the tracks of tidal streams with high precision. However, to motivate the work of this chapter, which was reported by Eyre & Binney (2009b) before the work of Chapter 5 was undertaken, we content ourselves with the examination of an N-body simulation of a stream superficially similar to the Orphan stream of Belokurov et al. (2007).

The full red curves in Fig. 2.2 show projections of an orbit (described in Table 2.1) outwardly similar to that underlying the Orphan Stream, from two viewing locations: the position of the Sun and a position 120° further round the Solar circle. Also shown in each projection are the locations of particles tidally stripped from a self-gravitating N-body model of a cluster launched on to that orbit. Clearly the particles provide a useful guide to the orbit of the cluster, but they do not precisely delineate it. Moreover, the relationship of the projected orbit to the stream depends on viewing angle. The line-of-sight velocities of the particles have a similar relationship to the orbit's line-of-sight velocity. Hence even with perfectly error-free observations, the track of a stream will not coincide with the progenitor orbit.

What would be the result of attempting to utilize the track of this stream in the procedure of B08? A set of points $\{l, b\}$ and $\{l, v_r\}$ was selected from the simulation output, by eye, to lie down the middle

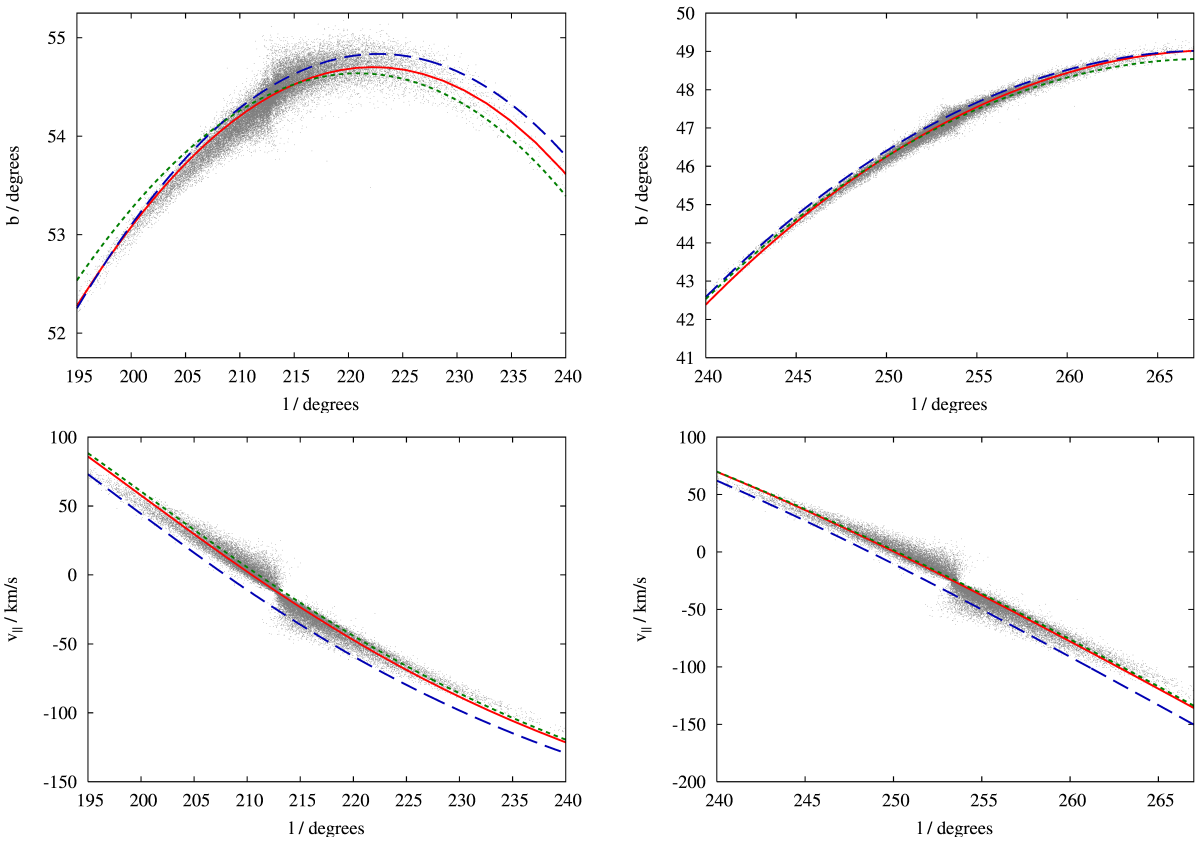

Figure 2.2: Full (red) lines: the orbit of a progenitor of an Orphan-like stream. Broken (green/blue) lines: orbits of a star now seen at either end of the tidal tail. Points: particles tidally stripped from an N-body model of the Orphan-like progenitor. Upper panels: distributions on the sky; lower panels: line-of-sight velocities. The N-body model had 60,000 particles set up as a King model with $W = 2$, $r_0 = 13.66\,\mathrm{pc}$ and $M_0 = 9381 M_\odot$, on the orbit detailed in Table 2.1, and evolved for 9.43 Gyr. The particles were advanced in time by the FVFPS tree code of Londrillo et al. (2003).

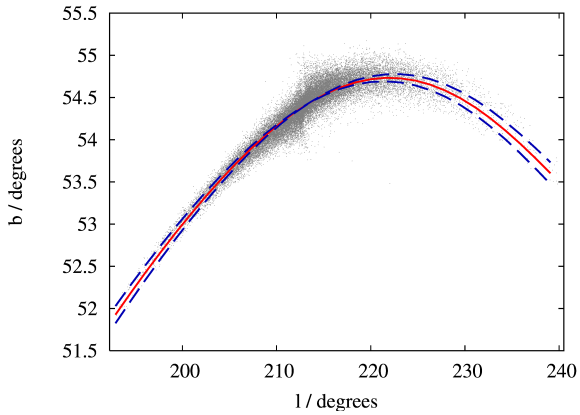

Figure 2.3: The dots are the projection of the simulated Orphan stream particles onto the sky. The red curve is a smooth polynomial fit to a set of points selected by eye to lie down the stream track. An input track $b(l)$ was defined by sampling 30 points from this curve, which was then used as input for the blue line in Fig. 2.1, and as baseline input for the data sets PD2–PD7. The blue dashed lines are the upper and lower bounds for PD2–PD7, as set by the penalty function p_{pos} (equation 2.23), and therefore represent the bow-tie region for these data sets.

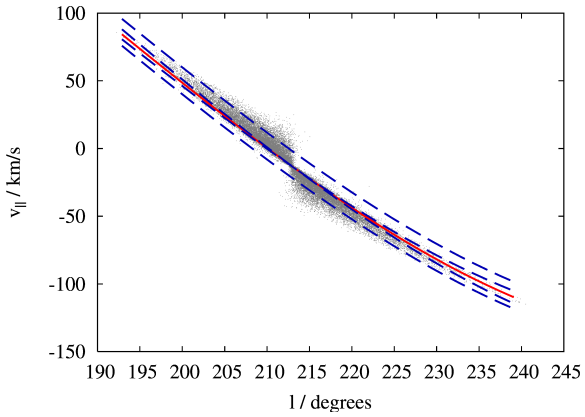

Figure 2.4: The dots are line-of-sight velocities for the particles of the simulated Orphan stream. The red curve is a smooth polynomial fit to a set of points selected by eye to lie down the stream track. An input track $v_r(l)$ was defined by sampling 30 points from this curve, and was used as input for the blue line in Fig. 2.1, and as baseline input for the data sets PD2–PD4. Also shown, as dashed blue lines, are the bow-tie regions for PD2 (narrow) and PD4 (wide). The bounds of these regions are enforced by the penalty function $p_{\rm vel}$ (equation 2.25).

of the stream. These sets were each fitted with a low-order polynomial to ensure smoothness, and these polynomials were sampled at 30 points to produce a set of input data, $b(l)$ and $v_r(l)$. The polynomial curves can be seen, along with the simulation data from which they derive, in Figs. 2.3 and 2.4.

The input data were used to solve the reconstruction equations (2.10) and (2.11) for a range of initial distances r_0. The blue curve in the left panel of Fig. 2.1 shows $\Delta E/E$ for the reconstructed trajectories. Unlike with the black curve, the minimum in the blue curve is very shallow, with energy conservation being half an order of magnitude worse than that of the erroneous input data, and $2\frac{1}{2}$ orders of magnitude worse than that of the perfect input data. It is not possible to make any statement about the true distance to the stream, or the validity of the potential, from the blue curve in Fig. 2.1.

Quite apart then from uncertainties in the observational data, the failure of tidal streams to properly delineate orbits has been shown to limit the diagnostic power of the B08 procedure. Therefore, in order to fully exploit the dynamical potential of streams, we have to understand how to infer the location of an underlying orbit from measurements of the stream.

For the moment we assume that any errors in measured quantities are negligible, which in practice means that they are small compared to the intrinsic width of the stream. This condition will certainly be satisfied by the on-sky coordinates. It will not always be satisfied by the velocity data, so we address the issue of velocity errors below.

2.3 Specifying stream tracks

With what precision could the track of an orbit be specified from the positions of the particles in Fig. 2.2? First we need to be clear that *any* orbit will do. Generally it will be convenient to use the orbit that passes through some point that lies near the centre of the observed stream, both on the sky and in line-of-sight velocity. In some circumstances this orbit will closely approximate the orbit of the centroid of the stream's progenitor, but there is no requirement for it to do so. Since a single fiducial point on the orbit can be chosen at will, this point is associated with vanishing uncertainty. As one moves away from this point, either up or down the stream, it becomes more uncertain where the chosen orbit lies, and the size of the uncertainties must increase. Hence the region of (l, b, v_r) space to which the orbit is confined by the observations is widest at its extremities and shrinks to a point at its centre.[2] We call this the "bow-tie region."

In the top panels of Fig. 2.2 the leading and trailing streams are not offset for much of the span, so our first guess may be that the orbit of the centroid runs right down the centre of the stream. In the lower panels the stream has a kink at the progenitor and our best guess is that the orbit through the point where the two halves of the tail touch runs near the lower edge of the left-hand half and near the upper edge of the right-hand half. In every case the uncertainty in the location of the orbit grows from zero at the centre to roughly half the width of the stream at its ends. Quantitatively, the largest uncertainty is then 0.15° for the top-left panel, $\sim 0.25^\circ$ for the top-right panel, and $\sim 5\,\mathrm{km\,s^{-1}}$ in the bottom panels.

2.4 Identifying dynamical orbits

B08 showed that given an orbit's projection onto the sky $[l(u), b(u)]$ and the corresponding line-of-sight velocities $v_r(u)$, the remaining phase space coordinates can be recovered by solving the differential equations (2.10) and (2.11).

If the input data used to solve those equations are not derived from an orbit in the same force field as is used to derive F_r, the reconstructed phase-space coordinates will not satisfy the equations of motion. B08 observed that violation of the equations of motion might cause the reconstructed solution to violate energy conservation, and therefore used rms energy variation down the track as a diagnostic for the quality of a solution. However, energy conservation is necessary but not sufficient to qualify a track as

[2] Observational error in v_r may lead us to demand that the orbit runs through a fiducial point that is not justified by the positional data. In this case, by making such a demand we may unfairly exclude from consideration the very orbit that we seek. In extreme cases, we may exclude all orbits, which would result in the erroneous rejection of the correct host potential. To avoid this, in practical use, the uncertainty associated with v_r for the orbit at the fiducial point will not be vanishing, but instead will be at a minimum and equal to the observational uncertainty on the v_r measurements themselves.

being an orbit. Here we construct a diagnostic quantity from residual errors in the equations of motion themselves, since orbits are defined to be solutions of these equations.

We first derive the equations of motion. In the Galactic rest-frame, the canonically conjugate momenta to the Galactic coordinates (r, b, l) are

$$\begin{aligned} p_r &= \dot{r}, \\ p_b &= r^2 \dot{b}, \\ p_l &= r^2 \cos^2 b \, \dot{l}, \end{aligned} \tag{2.13}$$

and the Hamiltonian is

$$\mathrm{H} = \frac{1}{2} p_r^2 + \frac{1}{2} \frac{p_b^2}{r^2} + \frac{1}{2} \frac{p_l^2}{r^2 \cos^2 b} + \Phi(r, b, l). \tag{2.14}$$

The equations of motion are therefore

$$\begin{aligned} \dot{p_r} &= \ddot{r} = \frac{p_b^2}{r^3} + \frac{p_l^2}{r^3 \cos^2 b} - \frac{\partial \Phi}{\partial r}, \\ \dot{p_b} &= \ddot{b} r^2 + 2 \dot{b} \dot{r} r = -\frac{p_l^2}{r^2} \frac{\sin b}{\cos^3 b} - \frac{\partial \Phi}{\partial b}, \\ \dot{p_l} &= r^2 \ddot{l} \cos^2 b - 2 r^2 \dot{l} \dot{b} \sin b \cos b + 2 r \dot{r} \dot{l} \cos^2 b = -\frac{\partial \Phi}{\partial l}. \end{aligned} \tag{2.15}$$

As in B08, when solving equations (2.10) and (2.11) we make extensive use of cubic-spline fits to the data. In the examples presented in B08 natural splines were used in order to avoid specifying the gradient of the data at its end points. Significantly improved numerical accuracy can be achieved by taking the trouble to specify these gradients explicitly. Given input data, we estimate the quantity $\mathrm{d}l/\mathrm{d}b$ at the end points by fitting a quadratic curve through the first three and last three points. $\mathrm{d}l/\mathrm{d}b$ is then computed at the location of the middle point of each set, and the very first and very last points are considered 'used' and thrown away. This quantity is then used in the geometric relation

$$\frac{\mathrm{d}u}{\mathrm{d}b} = \pm \sqrt{1 + \cos^2 b \left(\frac{\mathrm{d}l}{\mathrm{d}b} \right)^2}, \tag{2.16}$$

to compute $\mathrm{d}b/\mathrm{d}u$ at the ends of the track. The sign ambiguity is resolved by inspection of the directionality of the input data. We then use the geometric relation

$$\frac{\mathrm{d}l}{\mathrm{d}u} = \pm \sec b \sqrt{1 - \left(\frac{\mathrm{d}b}{\mathrm{d}u} \right)^2}, \tag{2.17}$$

to obtain $\mathrm{d}l/\mathrm{d}u$ at the ends of the track, where the sign ambiguity is resolved in the same way. We are now able to fit cubic splines through the input tracks, with the slopes at the ends of the $l(u)$ and $b(u)$ tracks given as above, but at this stage the track of $v_r(u)$ is fitted with a natural spline. The reconstruction equations (2.10) and (2.11) are now solved for $t(u)$, which is then fitted with a cubic spline, with the slopes at the ends given explicitly by equations (2.10) and (2.11) themselves.

We can now compute $l(t)$ and $b(t)$ and fit splines to them, with the slopes at the ends computed from $\mathrm{d}l/\mathrm{d}u$ and $\mathrm{d}b/\mathrm{d}u$ by the chain rule. The momenta (2.13) are now calculated explicitly, using the derivatives of the splines for $l(t)$ and $b(t)$ in place of $\dot{l}$ and $\dot{b}$. The slopes at the endpoints, $\mathrm{d}v_r/\mathrm{d}u$, can now be calculated from equation (2.15) and $\mathrm{d}t/\mathrm{d}u$; the $v_r(u)$ spline is refitted using these boundary conditions, the reconstruction repeated, and the momenta recalculated.

To compute a diagnostic quantity, the left- and right-hand sides of the equations of motion (2.15) are now evaluated explicitly. For each equation of motion we define a residual

$$R(t) = \dot{p}_{\mathrm{lhs}}(t) - \dot{p}_{\mathrm{rhs}}(t). \tag{2.18}$$

These residuals are used to compute, for each equation of motion, the diagnostic quantity

$$D = \log_{10}\left(\frac{\int_{t_1}^{t_2} \mathrm{d}t\, R(t)^2}{\int_{t_1}^{t_2} \mathrm{d}t\, \dot{p}_{\mathrm{lhs}}^2}\right), \tag{2.19}$$

where the residuals have been normalized by the mean-square acceleration and the times t_1 and t_2 correspond to the fifth and fifth-from-last input data points: numerical artefacts from the end regions, $0 < t < t_1$ and $t_2 < t < t_{\max}$, tend to dominate the integrated quantity and are not easily reduced by modifying the input; these regions are therefore excluded. The largest of the three values for D is used as the diagnostic quantity for that particular input.

2.4.1 Parameterizing tracks

Our strategy for identifying a stream's underlying orbit is to compute the diagnostic D (equation 2.19) for a large number of candidate tracks, and to find which candidates yield values of D consistent with their being dynamical orbits.

We start by specifying a baseline track across the sky $[l_{\mathrm{b}}(u'), b_{\mathrm{b}}(u')]$, where u' is a parameter that increases monotonically down the track from -1 to 1. Similarly, we specify associated baseline line-of-sight velocities $v_{r\mathrm{b}}(u')$. The baseline track is required to come closer to every data point than the given uncertainty at that point.

All candidate tracks should be smooth because orbits are. We satisfy this condition by expressing

the difference between the baseline track and a candidate track as a low-order polynomial in u'. For streams that cover a wide range of longitudes, the parameterization of candidate tracks is achieved by slightly changing the values of b and v_r associated with a given value of l from the values specified by the baseline functions. That is we write

$$\begin{aligned} b(u') &= b_{\rm b}(u') + \sum_{n=0}^{N} b_n T_n(u'), \\ v_r(u') &= v_{r\rm b}(u') + \sum_{n=0}^{N} a_n T_n(u'), \end{aligned} \tag{2.20}$$

where T_n is the $n^{\rm th}$-order Chebyshev polynomial of the first kind and a_n and b_n are free parameters. These $2N$ parameters are coordinates for the space of tracks that we have to search for orbits. When a stream does not stray far from the Galactic plane, candidate tracks are best parameterized by adjusting the baseline values of b and v_r at given longitude. In all examples in this chapter, the series in equations (2.20) are truncated after $N = 10$. A larger number of terms allows the correction function to produce tracks that represent orbits better, while using insufficient terms will fail to afford sufficient flexibility for the search procedure to find any orbits at all. However, using more terms makes the search procedure computationally more expensive, because the volume of the parameter space in which the search takes place grows exponentially with the number of terms, and locating solutions in this enlarged space requires commensurately more effort. The number of terms used is therefore a compromise between computational expense and the ultimate efficacy of the algorithm.

The space of tracks is defined by the a_n and b_n and one extra parameter, the distance to the stream, r_0, at the starting point $u = 0$ for the integration of equations (2.10) and (2.11).

We shall henceforth denote a point in the $(2N + 1)$-dimensional space of parameters by χ. Each χ is associated with a complete specification of all six phase-space coordinates for every point on the candidate orbit: l, b and v_r follow from the parameterization and the remaining coordinates are obtained by solution of the differential equations (2.10) and (2.11). Consequently, each χ corresponds to a value of D (equation 2.19) that quantifies the extent to which the phase-space coordinates deviate from a dynamical orbit in the given potential.

2.4.2 Searching parameter space

Dynamical orbits are found by minimizing the sum

$$D'(\chi) = D(\chi) + p(\chi), \tag{2.21}$$

where $p(\chi)$ is the sum of the penalty functions:

$$p(\chi) = \sum_i p_{i,\mathrm{pos}} + \sum_i p_{i,\mathrm{vel}} + p_s, \tag{2.22}$$

where

$$p_{i,\mathrm{pos}} = \begin{cases} \Delta_{i,\mathrm{pos}} & \text{if } \Delta_{i,\mathrm{pos}} > 1, \\ 0 & \text{otherwise,} \end{cases} \tag{2.23}$$

with

$$\Delta_{i,\mathrm{pos}} = \frac{|b(l_i) - b_\mathrm{b}(l_i)|}{\delta b(l_i)}. \tag{2.24}$$

Here δb_i is the width in b of the bow-tie region at l_i. Similarly

$$p_{i,\mathrm{vel}}(l) = \begin{cases} \Delta_{i,\mathrm{vel}} & \text{if } \Delta_{i,\mathrm{vel}} > 1, \\ 0 & \text{otherwise,} \end{cases} \tag{2.25}$$

with

$$\Delta_{i,\mathrm{vel}} = \frac{|v_r(l_i) - v_{r\mathrm{b}}(l_i)|}{\delta v_r(l_i)}. \tag{2.26}$$

Prior information about the distance to the stream is used by specifying the penalty function p_s to be

$$p_\mathrm{s} = \begin{cases} \beta\,|r_0 - r_{0\mathrm{b}}|\,/\delta r & \text{if } |r_0 - r_{0\mathrm{b}}| > \delta r, \\ 0 & \text{otherwise,} \end{cases} \tag{2.27}$$

where δr is the half-width of the allowed range in the distance r_0 to the starting point of the integrations and $r_{0\mathrm{b}}$ is the baseline value of r_0. These definitions are such that $p(\chi) = 0$ so long as the track lies within the region that is expected to contain the orbit, and rises to unity, or in the case of p_r to β, on the boundary and then increases continuously as the orbit leaves the expected region.

In practical cases the prior uncertainty in distance is large, and the obvious way to search for orbits is to set δr to the large value that reflects this uncertainty and then set the algorithm described below to work. It will find candidate orbits for certain distances. However, we shall see below that it is more instructive to search the range of possible distances by setting δr to a small value such as 0.5 kpc and searching for orbits at each of a grid of values of $r_{0\mathrm{b}}$. In this way we not only find possible orbits, but

we show that no acceptable orbits exist outside a certain range of distances. In this procedure the logic underlying δr is very different from that underlying δb and δv_r.

Since $p(\chi)$ is added to the logarithm of the rms errors in the equations of motion and since it increases by of order unity at the edge of the bow-tie region, the algorithm effectively confines its search to the bow-tie region, where $p = 0$. Thus at this stage we do not discriminate against orbits that graze the edge of the bow-tie region in favour of ones that run along its centre. Our focus at this stage is on determining for which distances dynamical orbits can be constructed that are compatible with the data. Once this has been established, distances that lie outside some range can be excluded from further consideration.

The space of candidate tracks χ is 21-dimensional, so an exhaustive search for minima of D' (equation 2.21) is impractical. Furthermore, the landscape specified by D' is complex. Some of this complexity is physical; the space should contain continua of related orbits, and ideally $D' \to -\infty$ at orbits. Hence deep trenches should criss-cross the space. Superimposed on this physical complexity is a level of numerical noise arising from algorithmic limitations in the computation of $D'(\chi)$. The limitations include the use of finite step sizes in the solution of equations (2.10) and (2.11) and the subsequent evaluation of $D'(\chi)$, the inability of a finite series of Chebyshev polynomials to precisely describe a true orbital track, as well as the difficulty in representing this series with a collection of sparse input points interpolated with splines. Reconstructed tracks are therefore never perfect orbital trajectories, even when the method is initially presented with error-free input data, and this is manifest as a non-zero minimum residual for each equation of motion. In practice, this minimum residual sets a lower limit on the returned values of $D'(\chi)$, which we refer to as the "numerical-noise floor".

On account of the complexity of the landscape that D' defines, "greedy" optimization methods, which typically follow the path of steepest descent, are not effective in locating minima. The task effectively becomes one of global minimization, which is a well studied problem in optimization.

We have used the variant of the Metropolis "simulated annealing" algorithm described in Press et al. (2002), which uses a modified form of the downhill simplex algorithm. In the standard simplex algorithm, the mean of the values of the objective function over the vertices decreases every time the simplex deforms. In the Press et al. algorithm the simplex has a non-vanishing probability of deforming to a configuration in which this mean is higher than before. Consequently, the simplex has a chance of crawling uphill out of a local minimum. The probability that the simplex crawls uphill is controlled by a "temperature" variable T: when T is large, uphill moves are likely, and they become vanishingly rare as $T \to 0$. During annealing the value of T is gradually lowered from an initially high value towards zero.

One vertex of the initial simplex is some point $\chi_{\rm guess}$, and the remaining $2N + 1$ vertices are obtained by incrementing each coordinate of $\chi_{\rm guess}$ in turn by a small amount. For the coefficients of the Chebyshev polynomial T_0 this increment is approximately the size of the allowed half widths, $\delta r, \delta v$ and

δb. Increments for coordinates representing coefficients of higher-order polynomials T_n are scaled as $1/n$. The overall size of these increments is therefore set by the size of the region within which we believe the global minimum to lie. It is important to note that in each generation of a simplex, the increments should independently have equal chance of being added to or subtracted from the values of $\chi_{\rm guess}$ so that no part of the parameter space is unfairly undersampled. The algorithm makes tens of thousands of deformations of the simplex while the temperature T is linearly reduced to zero. This entire process is repeated some tens of times, after which we have a sample of local minima that are all obtained from $\chi_{\rm guess}$.

We now update $\chi_{\rm guess}$ to the location of the lowest of the minima just found and initiate a new search. The entire process is repeated until the value of the diagnostic function $D'(\chi)$ hits a floor. When this floor lies higher than the numerical-noise floor, the attempt to find an orbit that is consistent with the assumed inputs has been a failure and we infer that no such orbit exists. When the floor coincides with the numerical-noise floor, we conclude that the corresponding χ specifies an orbit that is compatible with the inputs. An approximate value for the numerical-noise floor for a given problem may be obtained as follows: given input that perfectly delineates an orbit in the potential in use, the value of D' returned at the correct distance is approximately the numerical-noise floor. Conclusive proof that a candidate track with a particular value of D' is an orbit can be obtained by integrating the equations of motion from the position and velocity of any point on the track and ensuring that the time integration essentially recovers the track.

On account of the stochastic nature of the algorithm, an attempt to find a solution at a particular distance occasionally sticks at a higher value of D' than the underlying problem allows. This condition is identified by scatter in the values of D' reached on successive attempts and by inconsistency of these values with the values of D' achieved for nearby distances—we see from Fig. 2.5 that the function underlying the minima is smooth. When the magnitude of this scatter is significant, one can only confidently declare an attempt to find an orbit a failure if the D' achieved is consistently higher than the noise floor by more than the scatter; since the diagnostic measure D' quantifies the extent to which a candidate track satisfies the equations of motion, by definition, tracks with higher D' than the noise floor plus scatter cannot represent orbits.

When the observational constraints are weak, we expect several orbits to be compatible with them. In particular, we will be able to find acceptable orbits for a range of initial distances r_0. It is therefore important, for any given input, to run the algorithm starting from many different values of $r_{0\rm b}$ with δr set to prevent the algorithm straying far from the specified $r_{0\rm b}$. In this way, the full range of allowable distances can be mapped out, and dynamical orbits found for each distance in that range. In the case of significant scatter about the noise floor, the range of distances at which valid orbits are found is the range

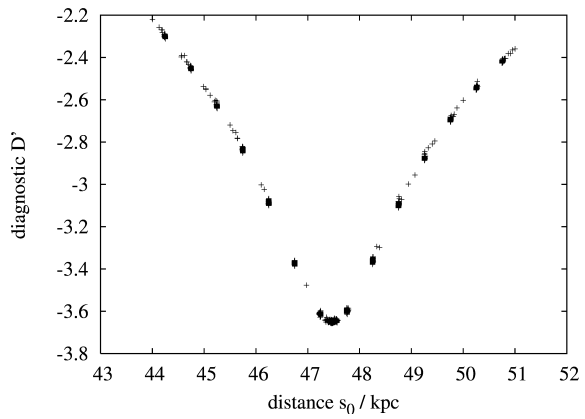

Figure 2.5: Values of the diagnostic function D' for candidate orbits reconstructed from the pseudo-data set PD1. Each group of crosses is associated with one of 15 ranges within which the starting distance r_0 was constrained to lie. For each such range the candidate orbit was reconstructed from 280 trial tracks, and each track yields a cross. In this figure the crosses are largely superimposed because the uncertainties are small and there is little scope for tweaking the track. This figure can be compared directly with the black curve of Fig. 2.1, the structure of which is clearly reflected in these results.

Table 2.2: Configurations of pseudo-data sets used to test the method.

	$\delta v_{r,\max}/\mathrm{km\,s^{-1}}$	$\delta v_{r,\min}/\mathrm{km\,s^{-1}}$	offset $v_r/\mathrm{km\,s^{-1}}$
PD1	2×10^{-3}	2×10^{-3}	0
PD2	4	0	0
PD3	6	2	0
PD4	10	10	0
PD5	6	2	2
PD6	15	10.5	10
PD7	4	0	0

within which solutions yield values of D' smaller than the noise floor plus scatter. Similar degeneracies in the parameters controlling the astrometry and line-of-sight velocities are less of a concern because if we have orbits that differ in these observables, we simply concentrate on the orbit that lies closest to the baseline track.

2.5 Testing the method

To test this method, we used the N-body approximation to the Orphan Stream shown in Fig. 2.2 as our raw data. The baseline input data $b_b(l)$ and $v_{rb}(l)$ were chosen to be the same 30-point samples of the smooth red curves from Figs. 2.3 and 2.4 as was used in §2.2.2 to test the B08 procedure. To the baseline data we attached uncertainties δb and δv_r, which through the penalty functions $p_{\rm pos}$ and $p_{\rm vel}$ (equations 2.23 and 2.25) constrain the tracks that the Metropolis algorithm can try. Details of the resulting pseudo-data sets are given below, and are summarised in Table 2.2.

In one case, PD1, the above baseline input data were replaced by those of a perfect orbit and δb and

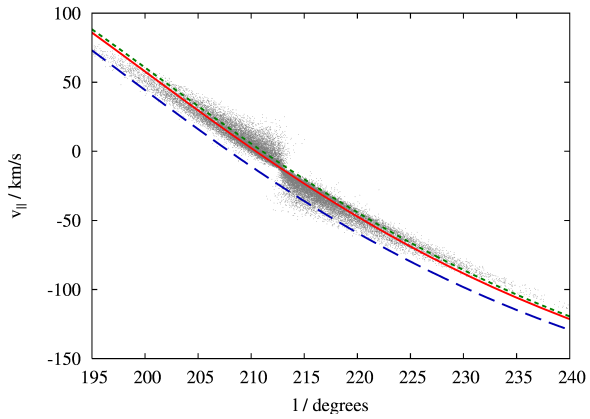

Figure 2.6: The broken lines show input tracks $v_r(l)$, as used for data sets PD5 (dotted green) and PD6 (dashed blue). The unmodified input of PD2 (full red curve) is shown for comparison. The dots are line-of-sight velocities for the particles of the simulated Orphan stream, from which the input was derived.

δv_r set very narrow (6 arcsec and $2 \times 10^{-3}\,\mathrm{km\,s^{-1}}$) in order to validate the reconstruction algorithm.

The uncertainty $\delta b(l)$ takes the same value for all the remaining pseudo-data sets because we assume that the astrometry is sufficiently precise for the uncertainty in position to be dominated by the offset of the stream from an orbit. In all cases, δb has a maximum value of 0.15° at the ends of the stream, falling linearly to zero at the position of the progenitor, consistent with the orbit of the progenitor seen in Fig. 2.2. Fig. 2.3 shows this input alongside the N-body data from which it was derived.

For the pseudo-data sets PD2 and PD3, δv_r is set to a maximum at the ends of the stream, and falls linearly to a minimum of zero and $2\,\mathrm{km\,s^{-1}}$ respectively, at the position of the progenitor. These examples represent a case (PD2) in which the uncertainty in radial velocity is dominated by stream width, and a case (PD3) in which there is a significant contribution from measurement error at a level that is easily obtainable with a spectrograph. The pseudo-data set PD4 represents the case in which the uncertainty in v_r is dominated by measurement errors: δv_r is held fixed at $10\,\mathrm{km\,s^{-1}}$, which is typical for the measurement errors in the line-of-sight velocities of SDSS stars in distant streams. Fig. 2.4 shows the input for these data sets.

For the pseudo-data sets PD5 and PD6, we added to the baseline data systematic offsets in $v_r(l)$ to mimic systematic errors in radial velocity. δv_r varies between a maximum and a minimum as in PD2 and PD3, with the values set to encompass the (assumed known) systematic bias. Fig. 2.6 shows the input for these data sets.

The pseudo-data set PD7 is identical to that of PD2, except that the number of raw (l, v_r) points was reduced to just three: one at either end of the N-body stream, and one at the location of the progenitor. A quadratic curve was perfectly fitted through these three points, and sampled at 30 locations to produce the baseline $v_r(l)$ input. δv_r is set to the same maximum value as PD2 at the outermost points; δv_r is

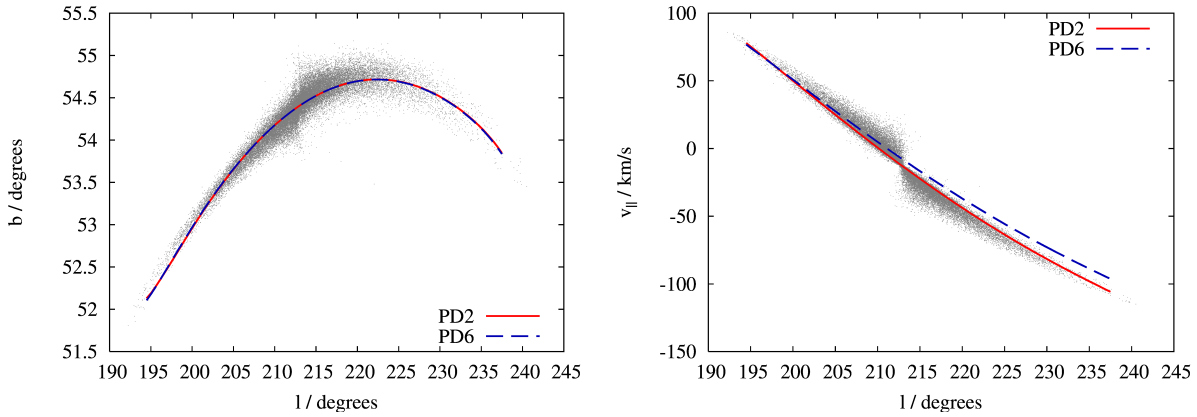

Figure 2.7: Left panel: projections onto the sky of two candidate orbits and the N-body data from which the input was derived. The full curves show the candidate orbit at 46 kpc from PD2, and the dashed curves show the candidate orbit at 43 kpc from PD6. Right panel: line-of-sight velocity down the track for the same candidate orbits. The penalty function (2.22) forces the sky projection and radial velocity curves of candidate orbits to be consistent with the data; the equivalent plots for all other tracks are very similar to these examples.

set to zero at the centre point. Only these three points are allowed to contribute to the penalty function (2.25) in this pseudo-data set.

In all of the examples, the penalty function (2.22) acts to constrain the candidate tracks to be consistent with the data. The contribution of the penalty function to D' is therefore zero in all examples. All candidate tracks are guaranteed to be consistent with the data, even if they do not represent dynamical orbits. Fig. 2.7 provides example plots of (l, b) and (l, v_r) for candidate tracks from PD2 and PD6 along with the raw N-body data from which the baseline data are derived. Equivalent plots for all candidate tracks are similar to these.

For PD1, PD2 and PD3, and each of 15 values of the baseline distance $r_{0\mathrm{b}}$, 280 optimization attempts were made, each involving Metropolis annealing for 24,000 simplex deformations, from an initial temperature of 0.5 dex. The starting distances were constrained by the penalty function p_s (equation 2.27) with $\beta = 10^6$ and $\delta r = 0.5\,\mathrm{kpc}$, so the Metropolis algorithm could only explore a narrow band in r_0. After 40 optimization attempts from a given starting point χ_{guess}, the starting point was updated to the end point of the most successful of these optimizations, and annealing recommenced from a high temperature. In total 6 of these updates were performed. With these parameters, a search at a single distance completes in 12 CPU-hours on 3GHz Xeon-class Intel hardware. PD4, PD5 and PD6 follow almost the same schema, except that 48,000 simplex deformations were made for each of 60 attempts at the same χ, which was updated 6 times. A single distance in the latter case took 36 CPU-hours to search: the computational load scales linearly with the number of deformations considered.

Fig. 2.5 shows the results obtained with PD1, which has very small error bars. The diagnostic

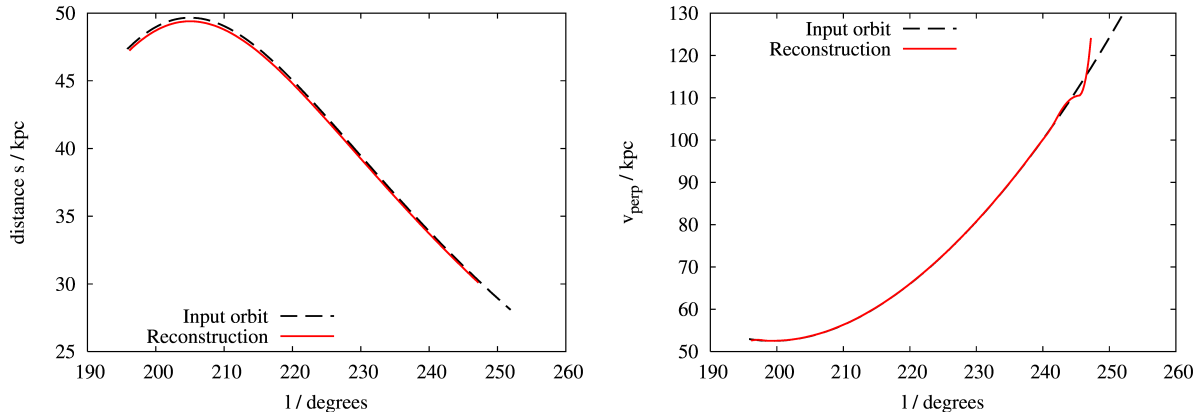

Figure 2.8: The left panel shows heliocentric distance, r, versus galactic longitude, l, for the best reconstruction from Fig. 2.5 and for the true orbit from which the input was generated. The two curves are close to overlying. The right panel shows tangential velocity, v_t, for the best reconstruction and for the true orbit: the departure of the reconstruction from the orbit near the endpoints is symptomatic of the problems near the endpoints that necessitate their excision from the diagnostic.

function D' has a smooth minimum that pinpoints the distance $r_0 = 47.4\,\mathrm{kpc}$ to the starting point with an uncertainty of $\sim 0.2\,\mathrm{kpc}$. The Metropolis optimization can significantly reduce D' by tweaking the input track only when r_0 is close to the truth. The depth of the minimum indicates the numerical-noise floor for this particular problem, and no significant scatter is seen between successive runs. We do expect the noise floor to be slightly different for PD2–PD6 because both the underlying orbits and the input are somewhat different. Fig. 2.8 shows that the reconstructed solution at the minimum overlies the input orbit almost exactly. We conclude that when the error bars are as small as in PD1, only one orbit is consistent with the data.

The left panels of Fig. 2.9 show the results obtained with PD2, PD3 and PD4. For PD2 and PD3, which have small to moderate error bars, the Metropolis algorithm reduces D' to the noise floor only for r_0 in the range $(44, 46)\,\mathrm{kpc}$. The scatter between runs is small, $\sigma_{D'} \sim 0.1$. The left panels of Figs. 2.10 and 2.11 show the reconstructed distances and tangential velocities associated with the best two solutions found at 44 and 46 kpc. With PD2, these reconstructions provide a distance estimate to the stream that is, at worst, 2 kpc in error, and a v_t estimate that is at worst $5\,\mathrm{km\,s^{-1}}$ in error. With PD3, the reconstruction is, at worst, 3 kpc and $10\,\mathrm{km\,s^{-1}}$ in error. For many sections of the orbits, the errors are less than stated. Thus the method can identify orbits consistent with the stream, and reject those that are inconsistent with it.

For PD4, which has large error bars comparable to those for velocity data from the SDSS, scatter between repeat runs is much larger, $\sigma_{D'} \sim 0.5$. Only distances $r_0 < 40\,\mathrm{kpc}$ can be confidently excluded, although a high-quality solution near $r_0 \sim 43\,\mathrm{kpc}$ provides a reconstruction in error by at most 2 kpc in distance and $5\,\mathrm{km\,s^{-1}}$ in v_t. Figs. 2.10 and 2.11 show us that ignoring this high-quality solution would

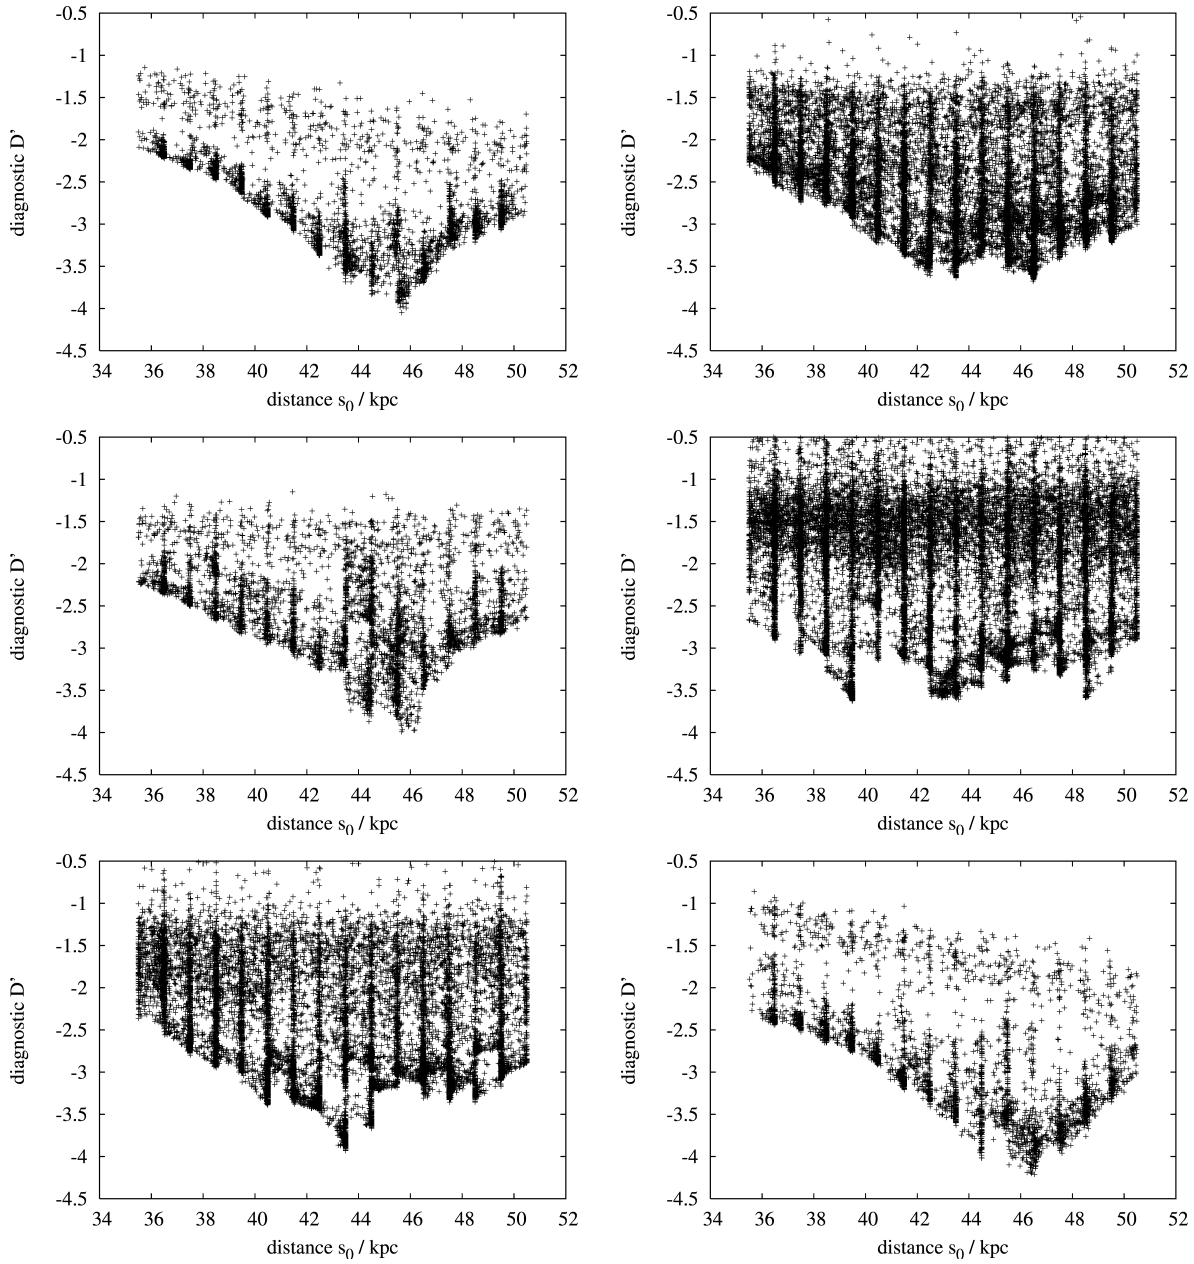

Figure 2.9: The same as Fig. 2.5 but for input data sets: PD2 (upper-left), PD3 (middle-left), PD4 (lower-left), PD5 (upper-right), PD6 (middle-right) and PD7 (lower-right). In PD2 and PD3, the noise floor $D' \sim -4$ is reached only for starting distances in the range $44 - 46$ kpc. In PD4 we find an isolated good solution at 43 kpc, but in contrast to what happens with data sets PD2 and PD3, as we change distance the value of D' oscillates around ~ -3.25 for most of the range, rather than varying smoothly. This behaviour arises because the volume of parameter space that has to be searched is large on account of the largeness of the velocity errors. Such an extensive volume of parameter space cannot be exhaustively searched with the allocated computational resource. Only orbits with $r_0 < 40$ kpc could be excluded with confidence. In PD5 and PD6, the noise floor $D \sim -3.5$ is now approached over a wider range in r_0: $42 - 47$ kpc in PD5 with some confidence, and $38 - 48$ kpc in PD6 with little confidence. As the errors increase, the search becomes a more arduous task, and patchy performance of this task is reflected in the rough bottoms to the graphs. In PD7, the results are very similar to those of PD2 and PD3, with a noise floor $D' \sim -4$ and solutions acceptable only in the range $44 - 46$ kpc.

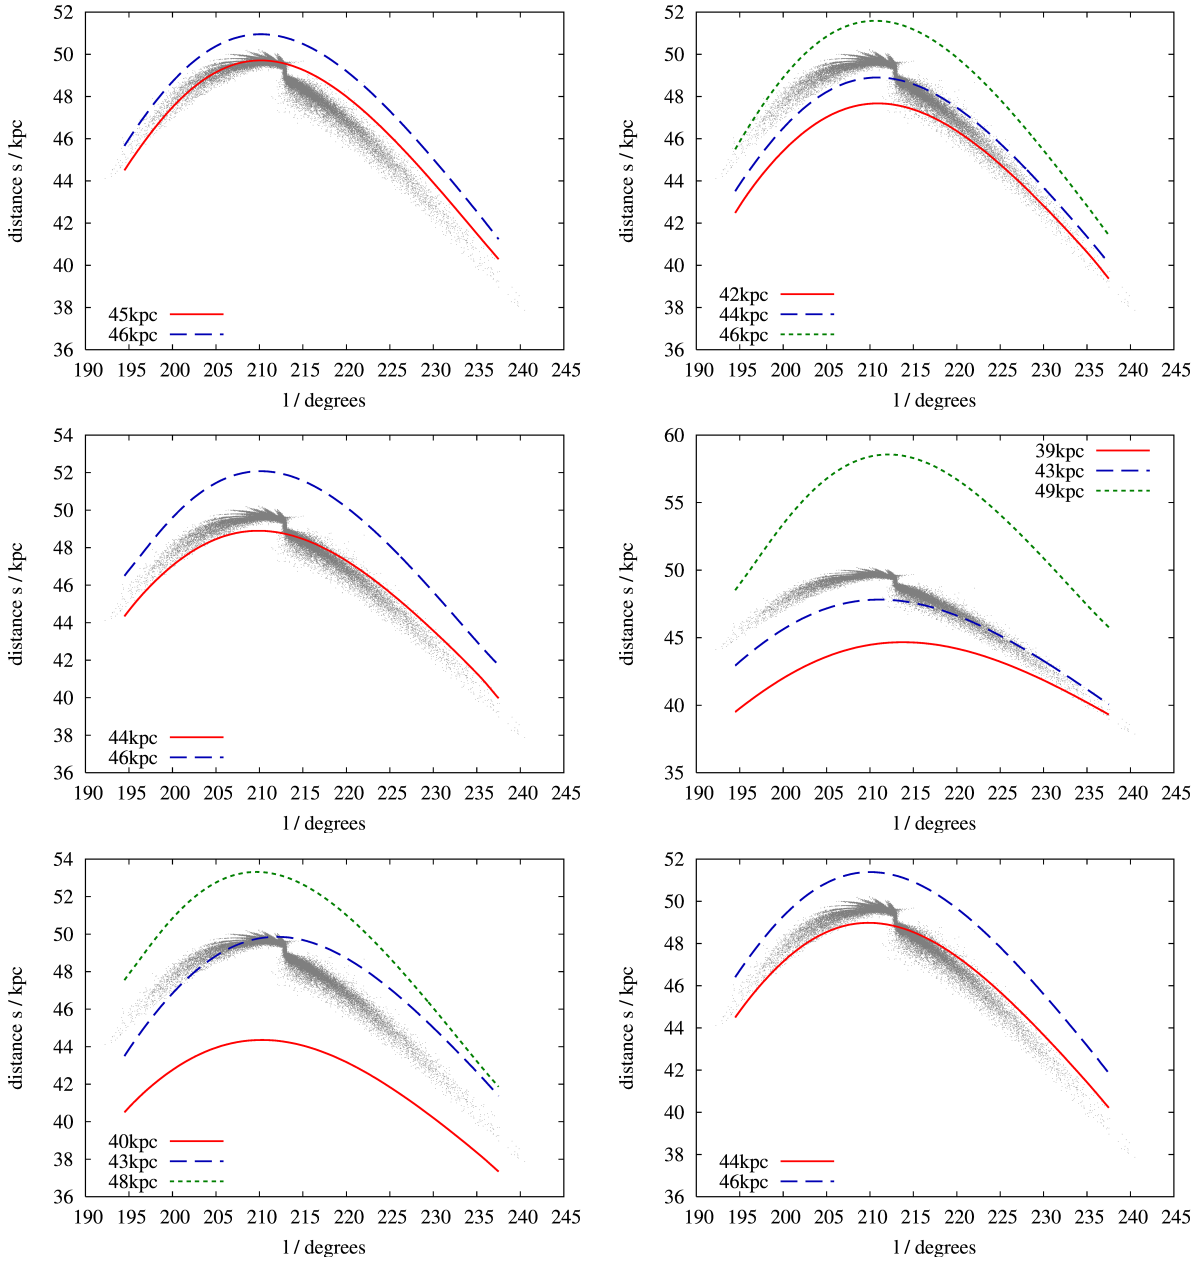

Figure 2.10: Heliocentric distance for selected reconstructed orbits from the data sets: PD2 (upper-left), PD3 (middle-left), PD4 (lower-left), PD5 (upper-right), PD6 (middle-right) and PD7 (lower-right). The tracks selected are those with lowest D' at distances for which D' approaches the noise floor in Fig. 2.9.

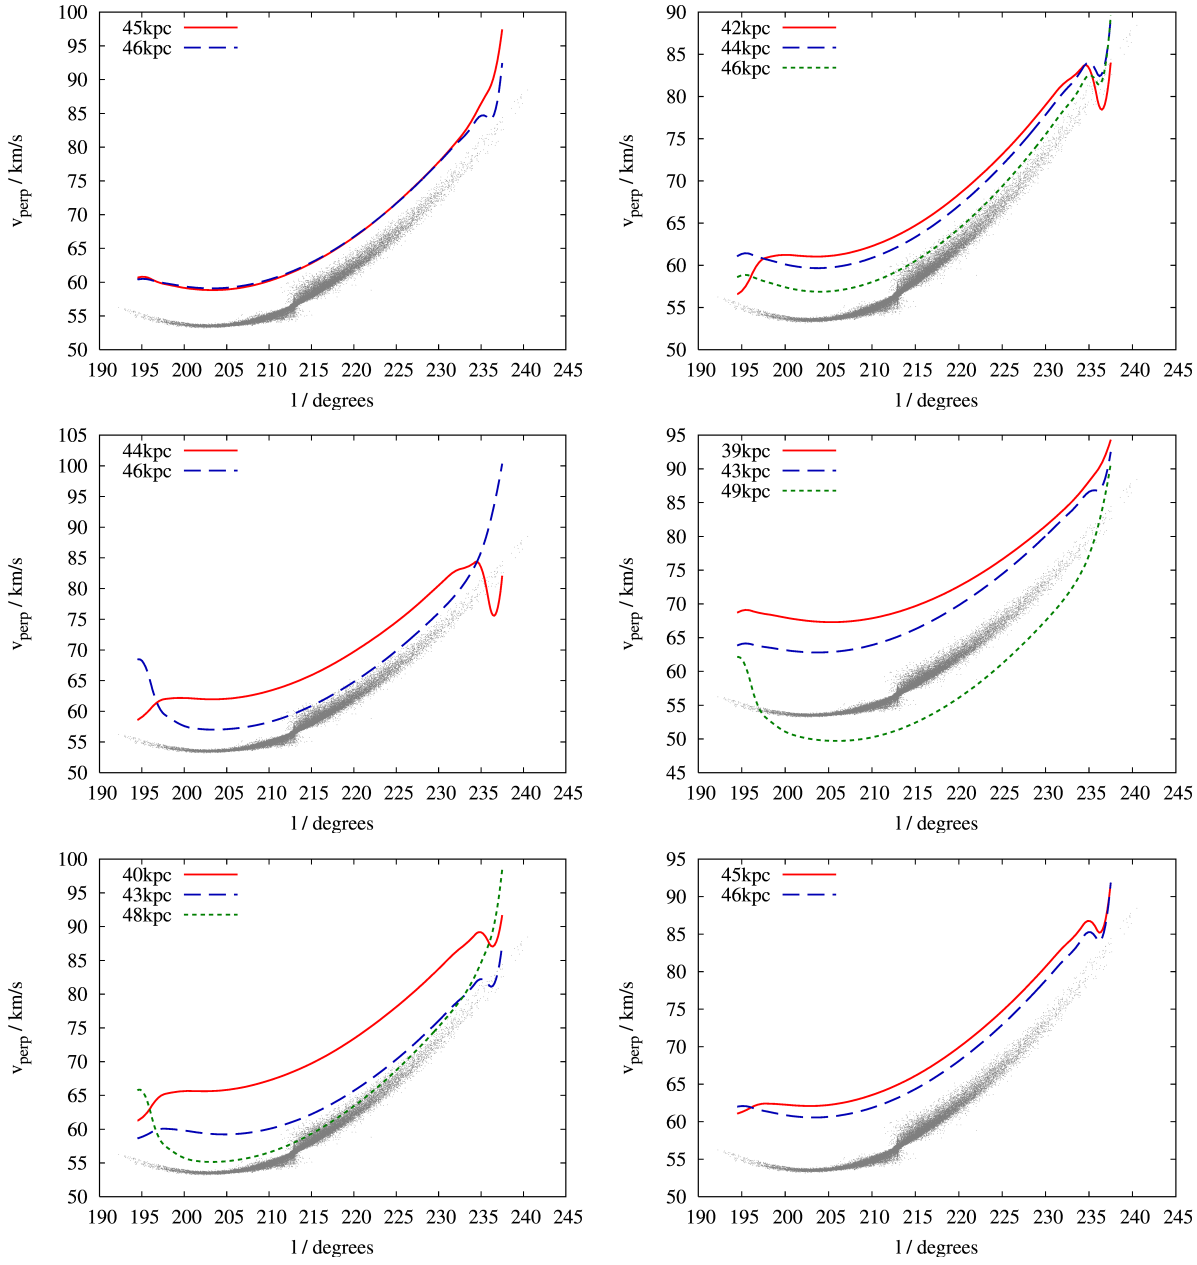

Figure 2.11: Tangential velocities for selected reconstructed orbits from the data sets: PD2 (upper-left), PD3 (middle-left), PD4 (lower-left), PD5 (upper-right), PD6 (middle-right) and PD7 (lower-right). The tracks selected are the same as in Fig. 2.10.

give reconstructed quantities in error by up to 5 kpc in distance and 12 km s^{-1} in v_t. Increasing the errors associated with the input clearly permits less satisfactory solutions to be returned, and also decreases the ability of the search procedure to find true orbits at particular distances, should they exist, by increasing the volume of parameter space it has to search. The former complaint is a physical statement about the limitations of the input data. The latter complaint is an algorithmic one, which may be remedied by providing the search with more computational cycles, or improving its efficiency.

The upper-right and middle-right panels of Fig. 2.9 shows the results obtained from input sets PD5 and PD6, in which offsets were applied to the input velocities. The reconstructed distances and tangential velocities for interesting tracks are shown in the upper-right and middle-right panels of Figs. 2.10 and 2.11. For PD5, the scatter between runs is $\sigma_{D'} \sim 0.5$, and the distance to the stream r_0 can be said to lie in the range $42 - 47$ kpc with some confidence. Fig. 2.10 shows that the stream does indeed lie in this range, so the algorithm has successfully corrected the small offset. The error in reconstructed distance and v_t would be 3 kpc and 8 km s^{-1} at worst.

PD6 demonstrates the limits of the method. On account of the large scatter in D', $\sigma_{D'} \sim 1$, we cannot identify the correct distance from the middle-right panel of Fig. 2.9. We expect the distance range for permitted orbits to be wide with this input, because the velocity error bars are large. PD5 and PD6 illustrate another problem of using large error bars: the search becomes harder because the parameter space to be searched is much larger. Consequently the plots of the results for PD5 and PD6 have rough bottoms, where the algorithm has failed to reach consistent minima for searches at adjacent distances. This can be remedied by re-running the search with more deformations and more iterations, and may be addressed in future upgrades to the search procedure.

PD7 demonstrates that the accuracy of the method is not necessarily significantly degraded when the number of velocity data points is substantially lower than the number of positional data points, as might be expected from a stream for which the only available radial velocities are those of giant stars. The results for PD7 are shown in the bottom-right panel of Fig. 2.9 and are directly comparable to those of PD2 and PD3 (the upper-left and middle-left panels). In particular, the range of allowed distances, $44 - 46$ kpc, is the same and the reconstructed orbits show comparable accuracy (cf. the upper-left and bottom-right panels of Figs. 2.10 and 2.11).

2.6 The effect of changing the potential

B08 demonstrated the ability of orbit reconstruction to diagnose the Galactic potential with astonishing precision when the track of an orbit on the sky is precisely known. Here we investigate the ability of orbit reconstruction to identify the correct potential when the track has to be inferred from realistic

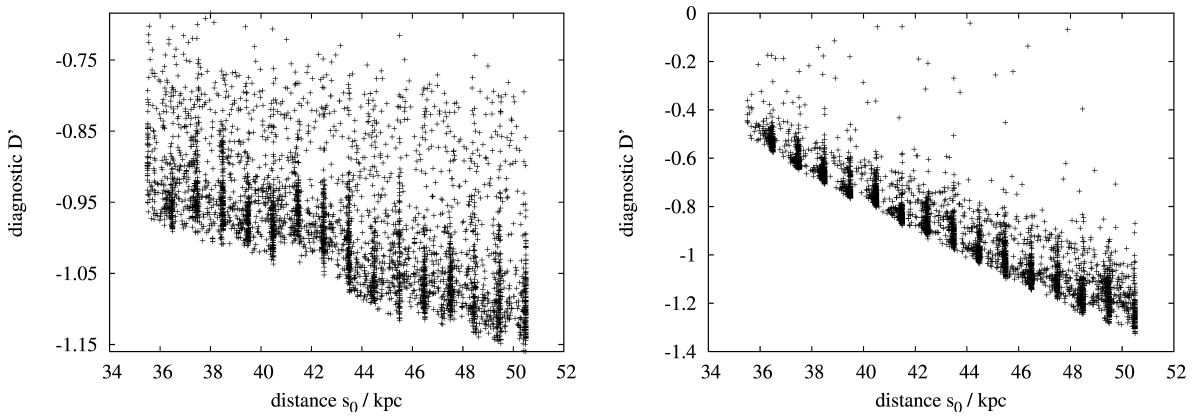

Figure 2.12: As Fig. 2.5, except the reconstruction (using PD2) takes place in a Kepler potential (left panel), and (right panel) a potential with $\Phi(r) \propto r$. In both cases, the potential parameters are set to generate approximately the same passage time along the stream as in Model II. Comparing with Fig. 2.9 shows the values of D' achieved are very poor, demonstrating that no orbits can be found in these potentials, which can therefore be excluded.

stream data.

We use as our input the PD2 data set from §2.5. The left panel of Fig. 2.12 shows the results of asking the algorithm to find orbits in a Kepler potential with mass $M = 4.18 \times 10^{11} M_\odot$, which produces roughly the same passage time along the stream as does Model II. The right panel shows the results obtained using a potential of the form $\Phi(r) = r f_r$, which gives a rotation curve of the form $v_c(r) = \sqrt{r f_r}$ with $f_r = 6.86 \times 10^2 (\,\mathrm{km\,s^{-1}})^2/\,\mathrm{kpc}$, again chosen to produce the same passage time along the stream as does Model II. These two potentials represent, respectively, relatively extreme falling and rising rotation curve models. We do not offer them as realistic candidates for the Galactic potential, but we intend to demonstrate that model potentials with approximately correct radial force, but incorrect shape, can be excluded using this method.

The distance range considered in both cases spans approximately ±15 percent of the true distance, which is the uncertainty one might expect in distances obtained by photometry. Since all values of D' are several orders of magnitude larger than the values one obtains with the correct potential, it is clear that no orbits can be found at the distances considered, and both potentials can be excluded.

Fig. 2.13 shows the results of reconstructing an orbit from the PD2 data set in a potential that differs from the Model II potential used to define the tidal stream only in that the halo mass has been varied by the specified ratio. This high-quality data set yields a sharp minimum in D' as a function of r_0 (cf. the upper-left panel of Fig. 2.9) and the left panel of Fig. 2.13 shows this minimum value of D' as a function of assumed halo mass. The minimum of D' lies at $D' < -4$ when the mass used lies within ~ 5 percent of the true value and $D' > -3.85$ otherwise. Hence, with high-quality data it is possible to constrain with remarkable precision the parameters of a model potential that is otherwise of the correct shape.

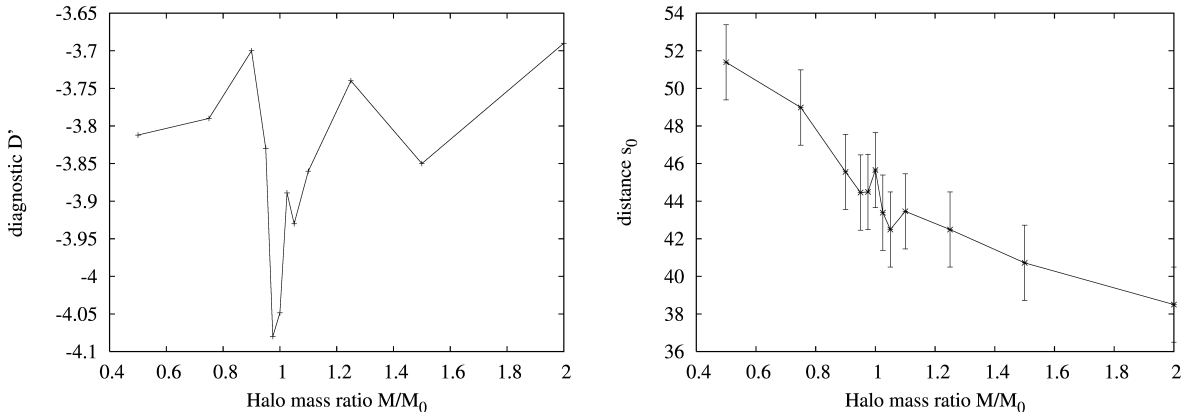

Figure 2.13: The left panel shows the minimum value of the diagnostic for attempted reconstruction of the PD2 input, versus the halo mass ratio of the Dehnen-Binney potential in which the reconstruction has been attempted. M_0 here is the value of the halo mass in the Model II potential. The right panel shows the characteristic distance, r_0, of the best solution, versus halo mass ratio. The error bars represent the approximate uncertainty in the recovered result.

Uncertainty estimates for potential parameters deduced in this way present some difficulty. Any parameters for which the minimum value of D' returned is greater than the noise floor are inconsistent with the data, and should be rejected. The uncertainty associated with each parameter is therefore that range over which solutions can be found that are consistent with the noise floor. Ambiguity arises, however, because the value of the noise floor itself must be estimated from the data, and is therefore uncertain. Furthermore, it is not clear with what probability an otherwise consistent set of parameters would return a value of D' higher than the noise floor, on purely stochastic grounds. In practice, potential parameters can only be confidently excluded if the scatter in returned value of D' is substantially smaller that the difference between D' and the noise floor, and yet the precise level of this confidence remains difficult to formally quantify.

The right panel of Fig. 2.13 shows the value of r_0 at which D' attains its minimum, as a function of assumed halo mass. We see that this value decreases systematically as the halo mass increases. Insight into this behaviour can be obtained by considering the discretized equation of angular-momentum conservation

$$\Delta(r\, v_t) = -r\Delta t\, \nabla_t\Phi = -r\frac{\Delta v_r}{F_r}\, \nabla_t\Phi, \tag{2.28}$$

where $-\nabla_t\Phi = (\hat{\mathbf{r}} \cdot \nabla\Phi\, \hat{\mathbf{r}} - \nabla\Phi)$ is the component of the force in the plane of the sky, and where Δ implies the change in a quantity between successive data points. By expanding the left side to first order in small quantities we can obtain an expression for Δv_t. Summing the changes in v_t along the track we

have

$$v_t(\text{end}) - v_t(\text{start}) = -\sum \left(\nabla_t \Phi \, \frac{\Delta v_r}{F_r} + v_t \frac{\Delta r}{r} \right). \tag{2.29}$$

An independent equation for v_t is

$$v_t = r\frac{\Delta u}{\Delta t} = rF_r \frac{\Delta u}{\Delta v_r}, \tag{2.30}$$

where we have used the radial component of Newton's second law. Equations (2.29) and (2.30) yield independent estimates of $v_t(\text{end}) - v_t(\text{start})$. The right side of equation (2.29) yields an estimate that is independent of the scaling of Φ but systematically decreases with increasing r, while the right side of equation (2.30) yields an estimate that is proportional to the scaling of Φ, but is almost independent of r, since $F_r \propto 1/r$. When the machine is asked to reconstruct the orbit with F_r taken too small, it can change the right side of equation (2.29) to match the new value of the right side of equation (2.30) by increasing r. In this way, the discrepancies between the left sides of equations (2.29) and (2.30), which contribute substantially to the diagnostic D', can be largely eliminated by increasing r as Φ is scaled down.

In principle this variation in reconstructed distance with the scaling of the potential could be combined with photometric distances to constrain the potential. Unfortunately, Fig. 2.13 shows that in the particular geometry under consideration, even a 10 percent distance error would produce a ~ 50 percent error in the estimate of the halo mass. Further work is required to discover what effect the geometry of a particular stream has on its sensitivity to the potential. Also of interest is whether simultaneously using multiple streams, or streams with multiple wraps (such as the Sagittarius Dwarf stream), can provide yet tighter constraints on the potential—B08 indicated that much could be gained from this approach, but more work is needed to extend the present schema to handle multiple independent segments of streams.

2.7 Conclusions

B08 demonstrated the tremendous diagnostic power that is available if one knows the track of an orbit on the sky and the associated line-of-sight velocities. Tidal streams are made up of stars that are on similar orbits. In particular, they roughly delineate the underlying orbit, but they do not do so exactly. We have presented a technique for identifying an underlying orbit and thus predicting the dynamical quantities that have not been observed, such as the distances to the stream and the proper motions of its particles.

A complete summary of our findings of this chapter is presented below. Following the summary is a discussion of the limitations of the work, accompanied by suggestions on how to overcome those limitations. Chapter 6 further reviews our findings in the context of contemporary astrophysics.

2.7.1 Summary

The technique presented involves defining a space of tracks on the sky and sequences of line-of-sight velocities that are consistent with the observational data, given the observational errors and the extent to which streams deviate from orbits. The equations of B08 are used to determine a candidate orbit for each track, and then the extent to which the candidate satisfies the equations of motion is quantified. This represents an enhancement over B08, in which only violations of energy conservation along a candidate were quantified. The resulting diagnostic quantity is used to search for tracks that could be projections of orbits in the Galactic potential. In practice the search is conducted for several possible distances to a fiducial point on the stream. If constraints on this distance are available from photometry, the computational effort of the search can be reduced by narrowing the range of distances for which searches need to be conducted. Our *a priori* assumptions are knowledge of the Galaxy's gravitational potential, and the Sun's velocity with respect to the Galactic centre.

The technique is moderately computationally expensive, in practice taking several CPU-core hours to process a given set of input for a single fiducial distance and potential combination. This computational expense arises from the large number of dimensions of the parameter space that the technique must search over: the parameter space must be sufficiently large to afford the technique the flexibility to derive a dynamical track from the input data, but the computational time that must be devoted to the search to ensure that the parameter space is appropriately explored grows exponentially with the dimensionality of the parameter space. The number of dimensions selected in practice is a compromise between these considerations. In mitigation, the search does lend itself to simple parallelization, and with a few dozen CPU-cores it is feasible to fully analyze a stream in a given potential in less than a day.

Our tests revolve around an N-body model of the Orphan Stream (Belokurov et al., 2007) on the assumption that it formed by the disruption of a globular cluster. We show that for this stream, which is $\sim 40^\circ$ long and 0.3° wide at its ends, distances to and tangential velocities of points on the stream can be recovered to within $\sim 2\,\mathrm{kpc}$ and $\sim 5\,\mathrm{km\,s^{-1}}$, respectively, if line-of-sight velocities accurate to $\sim 1\,\mathrm{km\,s^{-1}}$ are measured. As the errors in the measured radial velocities increase, the space of tracks that must be considered grows bigger and the search for acceptable orbits becomes more laborious. Moreover, the range of distances for which acceptable orbits can be found broadens. However, even with errors in radial velocities as large as $\pm 10\,\mathrm{km\,s^{-1}}$, the uncertainties in the recovered distances are no greater than ~ 10

percent and the recovered tangential velocities are accurate to better than ~ 20 percent. Zero-point errors in the input velocities that are reflected in appropriately wide error bars broaden the range of acceptable orbits but do not skew the results.

We have shown that the method maintains its accuracy even when very few radial velocity points are used to define the input, as might be necessary when radial velocities can only be measured for giant stars. In our tests, comparable results were obtained from pseudo-data based on only three velocity measurements and from pseudo-data based on fifteen velocity measurements. Indeed, the results obtained with three accurate velocity measurements were significantly superior to those obtained with fifteen lower-quality measurements. Naturally, exactly how many points are required to provide well-determined input will depend on the shape of the radial velocity curve along the stream in question.

We expect the accuracy of reconstructions to depend on the geometry of the problem in hand. In particular, we expect streams at apocentre, where families of orbits are compressed both on the sky and in radial velocity, to yield poorer results than streams away from apocentre. Unfortunately, streams are most likely to be discovered at apocentre because both trajectory compression and low proper motions around apocentre lead to a high density of stars at apocentre. We expect streams that are relatively narrow to produce more accurate results, because the permitted deviation of the orbit from the stream is then low. We also expect to have more difficulty reconstructing orbits from streams that contain a visible progenitor, since the potential of the progenitor will cause orbits in the progenitor's vicinity to differ materially from orbits in the Galaxy's underlying potential.

B08 suggested that it should be possible, if sufficiently accurate input is provided, to constrain the Galactic potential, since the wrong potential will not admit an acceptable orbit. We have tested this possibility for input with realistic errors. We find that two potentials of significantly different shape, the Kepler potential and $\Phi(r) \propto r$, are clearly excluded. We have also tested for changes in scaling of an otherwise correctly-shaped potential, by varying the mass of the assumed dark halo around the value used to make the pseudo-data. In this case, we find the correct potential is identified, with the diagnostic quantity generally worsening as the halo mass moves away from its correct value by more than ~ 5 percent. We further find a consistent relationship between the reported stream distance and the halo mass with which the reconstruction takes place. Although the reported distance is only weakly dependent upon halo mass, this does open the possibility of using alternative distance measurements, such as photometric distances, in conjunction with these techniques to constrain the Galactic potential.

Further work is necessary to determine a full scheme to recover parameters of the potential from stream data. Also in question is the extent to which simultaneous reconstruction of multiple streams, and reconstruction of streams with multiple wraps around the Galaxy, might provide stronger constraints on the Galactic potential than the short section of a single wrap that we have considered. It may also

prove possible to refine the other main assumption of our scheme, the location and velocity of the Sun.

2.7.2 Discussion

It is interesting to compare our method of finding orbits of progenitors with the traditional N-body method. Firstly, our method explores each orbit at a tiny fraction of the computational expense of N-body modelling, so it is feasible to automate the search of orbit space. Moreover, the search lends itself to easy parallelization. Finally, whereas only a successful attempt to model a stream with N body simulation yields an interesting conclusion, our method can show that no orbit is consistent with a given range of distances. This latter facility is particularly powerful, because by guaranteeing that no consistent orbit can be found, it is possible to test our assumed form for the Galactic potential.

The most serious limitation on the applicability of this work is the lack of high precision radial velocity measurements for streams. At the time of preparation of the Eyre & Binney (2009b) work on which this chapter draws, the radial velocity data available for a long, cold stream were scant: perhaps the best available were the two data points for the Orphan stream of Belokurov et al. (2007), which were held by the authors to be "suggestive rather than conclusive".

Since then, the situation has improved somewhat. Newberg et al. (2010) have recently reported radial velocity measurements for 7 fields along the Orphan stream, deriving from recently released SDSS/SEGUE spectra of giant stars. The measurements have associated uncertainties of $\sim 10\,\mathrm{km\,s^{-1}}$. The GD-1 (Grillmair & Dionatos, 2006) and Anticentre streams (Grillmair, 2006b) are closer to the Sun than is the Orphan stream, and therefore make easier targets for observations. Grillmair et al. (2008) obtained radial velocity data for two fields along the Anticentre stream (Grillmair, 2006b), with an accuracy of $\sim 7\,\mathrm{km\,s^{-1}}$. Koposov et al. (2010) combined SDSS/SEGUE spectra and purposely-obtained spectra of 24 main sequence and turn-off stars, and although they do not do so, their results could perhaps be organized into four or five fields with uncertainties of $\sim 4\,\mathrm{km\,s^{-1}}$ in each.

The procedures in this chapter could be immediately applied to all these streams. However, our experiments with pseudo-data indicate that the strongest constraints on the potential are to be had from $\sim 1\,\mathrm{km\,s^{-1}}$ precision data, rather than $\sim 10\,\mathrm{km\,s^{-1}}$ precision data. Except for perhaps the Koposov et al. (2010) observations of GD-1, the available radial velocity data do not meet this requirement.

Part of the problem stems from the need to use only the brightest stars in order to obtain useful spectra on the small-aperture telescopes that have so far been used for this work. The density of such stars associated with tidal streams is low, which are by their nature dynamically old structures, and are often barely distinguishable from the field stars even in the main sequence. The few giant stars that are observed are insufficient in number to beat down the observational random errors to the required level.

The problem is made worse by the difficultly in convincingly identifying giant stars with streams, with the small statistics involved leading to the possibility of a single contaminant star ruining an entire field.

Obtaining radial velocity measurements of streams with the required accuracy will necessitate the observation of main-sequence stars. Observations of main-sequence stars in distant streams such as the Orphan stream would require the use of 8 m-class equipment for what is superficially unglamourous work. However, the results of this chapter indicate that the scientific rewards of such an undertaking would be far-reaching.

Chapter 3

Fitting orbits to streams using proper motions

3.1 Introduction

The previous chapter has demonstrated the remarkable diagnostic power that orbit reconstruction from observed stream tracks can have, when coupled with radial velocity measurements down those streams.

However, it was proffered that perhaps the most serious problem afflicting such work is the difficulty in obtaining precise radial velocity measurements of stars in such streams, and particularly of the faint main-sequence stars that constitute the bulk of the prospective data.

One possible workaround to reconstruct phase-space tracks using measured proper motions rather than radial velocities. This is attractive because, although it is much harder to make an impromptu proper-motion measurement than it is to take a spectrum, it *is* possible to measure very many of the former simultaneously, down to the faint apparent magnitudes of main-sequence stars in distant streams.

This kind of work is already possible to some extent with the combined SDSS-USNO catalogue of Munn et al. (2004), as demonstrated by Koposov et al. (2010), who obtain comparatively high precision proper-motion measurements for the GD-1 stream (Grillmair & Dionatos, 2006) from lower-precision Munn et al. raw data, on account of the hundreds of main-sequence stars they are able to identify. By comparison, their spectroscopically-obtained radial velocity data number only in the dozens, with even this frugal number making GD-1 the best observed object in its class. Furthermore, with the advent of advanced astrometric projects such as the Pan-STARRS survey (Kaiser et al., 2002, currently commissioning), and its follow-on project LSST (Tyson, 2002), the future availability of large amounts of high-quality proper-motion data for main-sequence stars is assured. We note that Pan-STARRS proper

motions will be obtained for stars seven visual magnitudes fainter than those to be observed by Gaia (Perryman et al., 2001).

In the work of this chapter, much of which has been published in an article by Eyre & Binney (2009a), we derive a set of equations that can reconstruct full-phase space information from an on-sky track, if proper-motion magnitudes are measured everywhere along that track. The application of this technique is logically identical to the reconstruction of trajectories using radial-velocity data (Jin & Lynden-Bell, 2007; Binney, 2008; Chapter 2). We will demonstrate numerically that the reconstructed tracks can be tested for dynamicity, just as were those in Chapter 2. We further illustrate that the ability to use reconstructed tracks as probes of the Galactic potential is retained when using proper-motion data.

The remainder of this chapter is arranged as follows. §3.2 derives a set of reconstruction equations that use proper motions to generate candidate orbit trajectories. §3.3 demonstrates the efficacy of the reconstruction equations with a numerical test. §3.4 discusses our findings and sums up.

3.2 Reconstructing orbits with proper motions

As in Chapter 2, we derive a reconstruction algorithm on the basis that a stream precisely delineates an orbit. The probable effects of this incorrect assumption will be noted when applications of the work are discussed.

Consider two inertial frames of reference: one in which the Galactic centre is at rest (the grf), and another in which the Sun is instantaneously at rest (the hrf). In the grf, let $\mathbf{x}_0$ be the position vector of the Sun, $\mathbf{x}$ the location of a star in the stream and let $\mathbf{r}$ be the vector from the Sun to the star. By definition

$$\mathbf{r} = r\hat{\mathbf{r}} = \mathbf{x} - \mathbf{x}_0, \tag{3.1}$$

where $\hat{\mathbf{r}}$ is the direction from the Sun to the star. Now let $\mathbf{x}'$ be the vector obtained by transforming $\mathbf{x}$ from the grf frame to the hrf. Again, by definition

$$\mathbf{r}' = r'\hat{\mathbf{r}}' = \mathbf{x}' - \mathbf{x}'_0. \tag{3.2}$$

Consider the measurement of $\hat{\mathbf{r}}$ and $\hat{\mathbf{r}}'$, each made by one two observers, both at the location of the Sun, but each in his respective frame. Instantaneously the frames must coincide, so $\mathbf{r} = \mathbf{r}'$ at that given moment. However, the instantaneous derivatives $\mathrm{d}\hat{\mathbf{r}}/\mathrm{d}t$ and $\mathrm{d}\hat{\mathbf{r}}'/\mathrm{d}t$ will not be equal. The former derivative is tangent to the stream, and is the proper motion that would be made by an observer

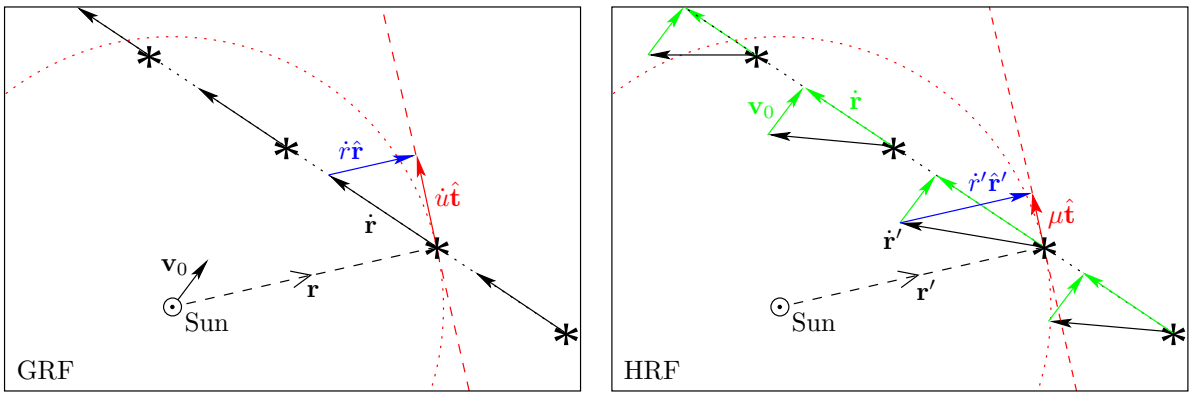

Figure 3.1: Diagrams illustrating the velocities of the Sun and of a stream of stars, as seen from (left) the Galactic rest-frame, and (right) the heliocentric rest-frame. Black arrows represent the velocities of objects. The components of velocity along the line-of-sight (blue arrows), and in the plane of the sky (red arrows), are shown for a star in the stream at position $\mathbf{r}$. A section of the plane of the sky at the radius r is shown as a dotted red circle, while the dashed red line shows the tangent to that circle at the position $\mathbf{r}$. Note that, even while the stream stars have been taken to follow one another in the grf, in the hrf the velocities of the stream stars do not point along the stream.

stationary with respect to the Galactic centre. The latter derivative is the terrestrial observable proper motion of the star, and includes a component perpendicular to the stream, due to the reflex motion of the Sun. Diagrams illustrating these configurations are shown in Fig. 3.1.

Now let u be the on-sky angle along the stream from some fiducial point. We may write

$$\frac{\mathrm{d}\hat{\mathbf{r}}'}{\mathrm{d}t} = \mu\hat{\mathbf{t}}; \qquad \frac{\mathrm{d}\hat{\mathbf{r}}}{\mathrm{d}t} = \dot{u}\hat{\mathbf{p}}, \tag{3.3}$$

where μ and $\hat{\mathbf{t}}$ are the magnitude and direction of the measured proper motion, while $\dot{u}$ and $\hat{\mathbf{p}}$ are the on-sky magnitude and direction of the star's motion along the stream. Space velocities measured in the grf differ from those measured in the hrf by the velocity of the Sun, that is

$$\frac{\mathrm{d}\mathbf{r}'}{\mathrm{d}t} = \frac{\mathrm{d}\mathbf{r}}{\mathrm{d}t} - \mathbf{v}_0, \tag{3.4}$$

where $\mathbf{v}_0$ is the grf velocity of the Sun. Differentiating equations (3.1) and (3.2), and combining the results with equation (3.3) and equation (3.4), we obtain

$$\begin{aligned}\frac{\mathrm{d}\mathbf{r}}{\mathrm{d}t} &= \dot{r}\hat{\mathbf{r}} + r\mu\hat{\mathbf{t}} \\ &= \frac{\mathrm{d}\mathbf{r}'}{\mathrm{d}t} - \mathbf{v}_0 = \dot{r}'\hat{\mathbf{r}}' + r'\dot{u}\mathbf{p} - \mathbf{v}_0 = \dot{r}'\hat{\mathbf{r}} + r\dot{u}\mathbf{p} - \mathbf{v}_0,\end{aligned} \tag{3.5}$$

where in the last step we have made use of the instantaneous equality $\mathbf{r} = \mathbf{r}'$. We note that $\mathrm{d}r'/\mathrm{d}t = \dot{r}'$ is the spectroscopically measured heliocentric velocity, and that $\mathrm{d}r/\mathrm{d}t = \dot{r} = v_r$ is the projection along

the line of sight of the star's velocity with respect to the Galactic centre. Equating the components in the plane of the sky, we have

$$r\mu\hat{\mathbf{t}} = r\dot{u}\hat{\mathbf{p}} - (\mathbf{v}_0 - \hat{\mathbf{r}}\cdot\mathbf{v}_0\,\hat{\mathbf{r}}) = r\dot{u}\hat{\mathbf{p}} - \mathbf{v}_s, \tag{3.6}$$

where $\mathbf{v}_s$ is the component of the Sun's velocity perpendicular to the line of sight. This equation has just two unknowns, $\dot{u}$ and r, and we can in principle solve for both through

$$\dot{u}\hat{\mathbf{p}} = \mu\hat{\mathbf{t}} + \frac{\mathbf{v}_s}{r}. \tag{3.7}$$

Specifically, since both $\hat{\mathbf{t}}$ and $\hat{\mathbf{p}}$ can in principle be deduced from the observations and $\mathbf{v}_s$ may be presumed known, we could determine r such that the right side is parallel to $\hat{\mathbf{p}}$, and then read off $\dot{u}$ from the magnitude of the right side. Startlingly, this permits the distance to the stream to be recovered from the observables with no additional assumptions. This method embodies the idea of Galactic parallax, and is examined further in Chapter 4. For now, we will consider that case where the uncertainty in the direction $\hat{\mathbf{t}}$ is significant, prohibiting the use of the Galactic parallax procedure. We proceed to eliminate $\hat{\mathbf{t}}$ by squaring up

$$(\dot{u})^2 - 2\frac{\mathbf{v}_s\cdot\hat{\mathbf{p}}}{r}\dot{u} + \frac{|\mathbf{v}_s|^2}{r^2} - \mu^2 = 0. \tag{3.8}$$

The roots of this quadratic equation in $(r\dot{u})$ are given by,

$$r\dot{u} = \mathbf{v}_s\cdot\hat{\mathbf{p}} \pm \sqrt{(\mathbf{v}_s\cdot\hat{\mathbf{p}})^2 + (r\mu)^2 - |\mathbf{v}_s|^2}. \tag{3.9}$$

The sign ambiguity is resolved as follows. Using this expression to eliminate $\dot{u}$ from equation (3.7), we find

$$\mu\,\hat{\mathbf{t}}\cdot\hat{\mathbf{p}} = \pm\sqrt{(\mathbf{v}_s\cdot\hat{\mathbf{p}})^2 - |\mathbf{v}_s|^2 + (r\mu)^2}. \tag{3.10}$$

Since μ is inherently positive, we see that the sign of the radical in equation (3.9) must be chosen to agree with the sign of $\hat{\mathbf{t}}\cdot\hat{\mathbf{p}}$. Even though the directions of individual proper motions may be uncertain, it should be possible to decide whether they are on average opposed to the direction of travel along the stream. This ambiguity resolved, equation (3.9) now makes $\dot{u}$ into a function of both r and quantities that can be determined from the observations.

Now let $\mathbf{F}(\mathbf{r}) = -\nabla\Phi(\mathbf{r})$ be the Galaxy's gravitational acceleration. We recall that when we resolve

the star's equation of motion along the line of sight, we obtain (equation 2.6)

$$\frac{\mathrm{d}v_r}{\mathrm{d}t} = \dot{u}\frac{\mathrm{d}v_r}{\mathrm{d}u} = F_r + \frac{v_t^2}{r}, \tag{3.11}$$

where $F_r = \hat{\mathbf{r}} \cdot \mathbf{F}$ is the radial component of the acceleration, $v_t = \sqrt{v^2 - v_r^2}$ is the magnitude of the on-sky component of the star's velocity, and where we have explicitly written the time derivative in terms of u to make it clear that we refer to phase along the stream, and not a measured flux.

Since equation (3.9) makes $\dot{u}$ a known function of r, in conjunction with equation (3.11) we can now write down a system of three coupled nonlinear ODEs for the unknowns along the stream

$$\begin{aligned}
\frac{\mathrm{d}r}{\mathrm{d}u} &= \frac{v_r}{\dot{u}}, \\
\frac{\mathrm{d}v_r}{\mathrm{d}u} &= \frac{F_r + r\dot{u}^2}{\dot{u}}, \\
\frac{\mathrm{d}t}{\mathrm{d}u} &= \frac{1}{\dot{u}}.
\end{aligned} \tag{3.12}$$

If initial conditions on (r, v_r, t) were given at some point on the stream, by integrating these equations we can obtain complete phase-space coordinates at every point along the stream. We can trivially set $t = 0$ at a fiducial point along the stream, and like in Chapter 2, we can guess an initial distance $r = r_0$ at that point. To obtain the initial condition on v_r we write

$$\begin{aligned}
F_t &= \frac{\mathrm{d}\mathbf{v}}{\mathrm{d}t} \cdot \hat{\mathbf{p}} = \frac{\mathrm{d}}{\mathrm{d}t}(v_r\hat{\mathbf{r}} + r\dot{u}\hat{\mathbf{p}}) \cdot \hat{\mathbf{p}} \\
&= v_r\frac{\mathrm{d}\hat{\mathbf{r}}}{\mathrm{d}t} \cdot \hat{\mathbf{p}} + v_r\dot{u} + r\ddot{u} \\
&= 2v_\parallel\dot{u} + s\ddot{u},
\end{aligned} \tag{3.13}$$

where we have used equation (3.3) to eliminate $\mathrm{d}\hat{\mathbf{r}}/\mathrm{d}t$. Hence

$$\frac{\mathrm{d}(r\dot{u})}{\mathrm{d}t} = \dot{u}\frac{\mathrm{d}(r\dot{u})}{\mathrm{d}u} = F_t - v_r\dot{u}. \tag{3.14}$$

The lhs of this equation is obtained by explicitly differentiating equation (3.9)

$$\frac{\mathrm{d}(r\dot{u})}{\mathrm{d}u} = \alpha + \frac{1}{\beta}\left(\alpha\gamma + \mu r^2\frac{\mathrm{d}\mu}{\mathrm{d}u} + \frac{r\mu^2 v_r}{\dot{u}} - \mathbf{v}_s \cdot \frac{\mathrm{d}\mathbf{v}_s}{\mathrm{d}u}\right), \tag{3.15}$$

where we have defined

$$\begin{aligned}\alpha =& \frac{\mathrm{d}\mathbf{v}_s}{\mathrm{d}u} \cdot \hat{\mathbf{p}} + \mathbf{v}_s \cdot \frac{\mathrm{d}\hat{\mathbf{p}}}{\mathrm{d}u}, \\ \beta =& \sqrt{(\mathbf{v}_s \cdot \hat{\mathbf{p}})^2 + (r\mu)^2 - v_s^2}, \\ \gamma =& \mathbf{v}_s \cdot \hat{\mathbf{p}}.\end{aligned} \tag{3.16}$$

We can evaluate the terms in these quantities as follows. By differentiation of $\mathbf{v}_s = \mathbf{v}_0 - \hat{\mathbf{r}} \cdot \mathbf{v}_0 \, \hat{\mathbf{r}}$, we find that

$$\begin{aligned}\frac{\mathrm{d}\mathbf{v}_s}{\mathrm{d}u} &= -\hat{\mathbf{r}} \cdot \mathbf{v}_0 \, \frac{\mathrm{d}\hat{\mathbf{r}}}{\mathrm{d}u} - \frac{\mathrm{d}\hat{\mathbf{r}}}{\mathrm{d}u} \cdot \mathbf{v}_0 \, \hat{\mathbf{r}}. \\ &= -\hat{\mathbf{r}} \cdot \mathbf{v}_0 \, \hat{\mathbf{p}} - \hat{\mathbf{p}} \cdot \mathbf{v}_0 \, \hat{\mathbf{r}},\end{aligned} \tag{3.17}$$

and so

$$\frac{\mathrm{d}\mathbf{v}_s}{\mathrm{d}u} \cdot \hat{\mathbf{p}} = -\hat{\mathbf{r}} \cdot \mathbf{v}_0, \tag{3.18}$$

and

$$\mathbf{v}_s \cdot \frac{\mathrm{d}\mathbf{v}_s}{\mathrm{d}u} = -(\hat{\mathbf{r}} \cdot \mathbf{v}_0) \, \mathbf{v}_s \cdot \hat{\mathbf{p}}, \tag{3.19}$$

all of which can be evaluated given the initial condition r_0 and the observables. Since equation (3.15) is linear in v_r, we can combine it with equation (3.14) and rearrange to give

$$v_r = \frac{1}{\beta + s\mu^2/\dot{u}} \left(\frac{\beta F_t}{\dot{u}} - \alpha\beta - \alpha\gamma + \mathbf{v}_s \cdot \frac{\mathrm{d}\mathbf{v}_s}{\mathrm{d}u} - \mu \frac{\mathrm{d}\mu}{\mathrm{d}u} s^2 \right). \tag{3.20}$$

Once an initial distance $r = r_0$ has been chosen, the right side of this equation can be evaluated from the observables at the fiducial point on the stream. The initial conditions required for the integration of equations (3.12) are then known, and full phase-space information for the stream is defined.

In practice, it is clear that not all values of r_0 will provide solutions to equation (3.7). In particular, if r is significantly too small, equation (3.7) will become insoluble at some point along the track. This pathology manifests itself as $\dot{u} \to 0$ which halts the solution of the differential equations (3.12). This event signals that the measured value of μ is too small to be consistent with the reflex motion of the Sun at the proposed distance, and it effectively sets a geometrical lower bound on r_0. There is no similar geometrical upper bound, since $\dot{u} \to \mu$ as $r_0 \to \infty$ and equation (3.7) always has a solution. In the range above the lower bound, r_0 must be searched over, with dynamics used to isolate any physical solutions,

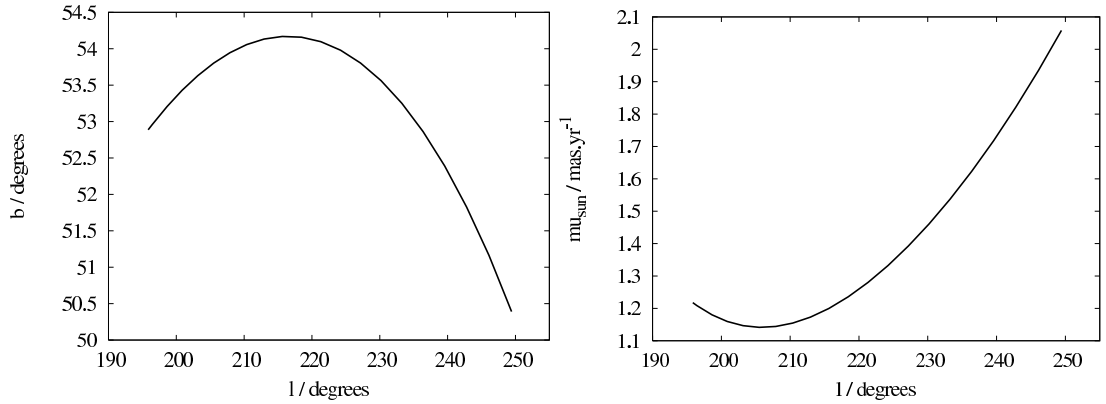

Figure 3.2: Test input for the algorithm. The input is derived from a perfect orbit in the Model II potential, and is described in the text. Left panel: on-sky projection of the orbit. Right panel: heliocentric proper-motion magnitude.

just as for the radial-velocity reconstruction in Chapter 2.

3.3 Tests

The implementation of the radial-velocity reconstruction algorithm from Chapter 2 was adapted to solve equations (3.12). Given an on-sky track $[l(u), b(u)]$ and heliocentric proper motions $\mu(u)$, along with an assumed potential $\Phi(\mathbf{r})$ and an initial distance r_0, the algorithm returns a trajectory for which full phase-space information is known. Aside from the system of equations being solved, the sole difference between the implementation in Chapter 2 and that here is the following minor procedural upgrade: instead of the fiducial point from which the integration begins being located at one of the ends of the input track, it was instead located in the middle. This leads to a somewhat more accurate start to the integration, since the data constrain the derivatives much more tightly in the middle of the track than they do at the ends.

To test the algorithm, a segment of the PD1 Test Orbit from Chapter 2 was integrated in the Model II potential of Binney & Tremaine (2008). This segment, which is identical to that used to test the radial velocity reconstruction algorithm in Fig. 2.1, was used to generate the input data sets $[l(u), b(u)]$ and $\mu(u)$, assuming the Sun to be on a circular orbit of radius $R_0 = 8\,\mathrm{kpc}$, at which Model II generates a circular velocity $v_c = v_0 = 227\,\mathrm{km\,s^{-1}}$. The input data sets, which represent a perfect orbit in the Model II potential, are shown in Fig. 3.2.

The reconstruction equations (3.12) were then solved for this input, and for a range of initial distances r_0, with each initial distance producing a different reconstructed trajectory. As in Binney (2008) and §2.2, the trajectories were tested for dynamicity by examining the conservation of orbital energy along

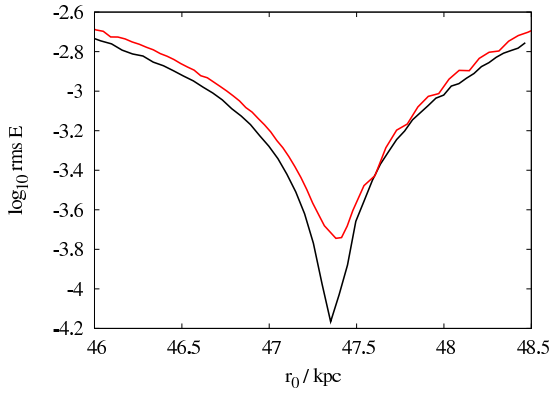

Figure 3.3: Diagnostic quantity $\Delta E/E$ for a range of initial distances, r_0, for input corresponding to a perfect orbit in the Model II potential. Black line: the procedure performed in the Model II potential. Red line: the procedure performed in a potential with halo asymmetry parameter $q = 1$, but otherwise identical to the Model II potential.

them. The normalized rms orbital energy variation

$$\frac{\Delta E}{E} \equiv \sqrt{\frac{\langle E^2 \rangle}{\langle E \rangle^2} - 1}, \tag{3.21}$$

was computed for each trajectory, where $E_i = v_i^2/2 + \Phi(\mathbf{r}_i)$ at a point $\mathbf{r}_i$ along the track.

The black curve of Fig. 3.3 shows the diagnostic quantity $\Delta E/E$ as a function of initial distance r_0, where the reconstruction algorithm has been given the correct form of the host potential $\Phi(\mathbf{r})$. As with the reconstruction using radial velocities in Chapter 2, the procedure has highlighted the correct distance to the stream with superb precision. Indeed, in this example, the energy conservation of the best reconstructed trajectory exceeds, by half an order of magnitude, the energy conservation of the equivalent track when reconstructed from radial velocity data. The reasons for this are twofold. Firstly, the upgraded technique of beginning the integration from the centre of the track has reduced numerical noise slightly. Secondly, the two reconstruction techniques, while intellectually equivalent, are mathematically distinct, and it is understandable that one technique may fare better than another, given otherwise identical input.

The red curve in Fig. 3.3 is for reconstructions using the same input data as the black curve, but in this case we have asked the algorithm to reconstruct in a potential similar to Model II, but for which the halo asymmetry parameter (Binney & Tremaine, 2008, Table 2.3) has been changed from its Model II value of $q = 0.8$ to $q = 1.0$. Thus, the algorithm is attempting to reconstruct trajectories using a potential which is less flattened than the truth.

A minimum in the red curve is still very distinct, and the minimum indeed occurs at approximately the correct distance, as can be seen by comparison with the black curve. We understand this, because we have not altered the halo mass, only its shape. Thus, the magnitude of the force at the distance of

the stream remains approximately correct. The results of §2.6 showed that the force strongly determines the distance r_0 at which the optimum trajectory is found. However, the minimum in the red curve shows substantially worse energy conservation than does the black curve, being half an order of magnitude above the noise floor, as determined by the minimum of the black curve. Thus, as in a reconstruction using radial velocities, when given perfect input, the algorithm can successfully diagnose an incorrect potential.

3.4 Conclusions

We have complemented the work of Binney (2008) and Chapter 2 by showing that, given proper motions measured everywhere along a segment of a single orbit, then full phase-space information can be reconstructed for that orbit, just as it can if radial-velocity measurements are known.

However, proper motions are two-dimensional vectors rather than scalars like line-of-sight velocities, and this fact is rather useful. We have found that, in principle, an orbit can be recovered from proper-motion measurements on purely geometric grounds, with no assumptions made about the potential, and no scanning over distances to find dynamical tracks required. Since any orbit so recovered *must* correspond to an orbit in the correct potential, the latter would be tightly constrained.

We explore this exciting discovery is full detail in Chapter 4. In the work of this chapter, we have chosen to sacrifice some of the diagnostic power of the proper motions by utilizing only their magnitude, and not their direction, which we deem to be the most difficult quantity to measure. The resulting scheme is logically identical to the scheme of Binney (2008) for reconstructing trajectories from radial-velocity data. Just as in the latter, trajectories reconstructed from proper motions are not forced to be dynamical, and we must hunt for those that are, by computing a diagnostic criterion for each candidate trajectory.

Our tests have indicated that reconstructions using proper motion data are at least as potent as are radial-velocity reconstructions when it comes to both isolating dynamical orbits, and diagnosing the potential. However, like with the radial velocity reconstruction, we expect random input errors, and the failure of actual stream tracks to delineate individual orbits, to be destructive of our ability to identify dynamical solutions.

In order to unlock the diagnostic power of streams using proper-motion measurement, it is likely that techniques similar to those presented in Chapter 2 will be necessary, affording the algorithm the ability to detect those dynamical solutions that can be found with constrained modification of its input. In principle, the adaptation of the procedures to the new algorithm ought to be straightforward, given the intellectual similarity between the two sets of reconstruction equations, and given the similarity in our

method of implementing them.

However, we will leave this interesting endeavour for future work. Instead, in the following chapter we will revisit our assertion that recovery of the distance to a stream star can proceed from nothing more than measurement of its on-sky coordinates and proper-motion vector.

Chapter 4

Galactic parallax for the tidal stream GD-1

4.1 Introduction

Conventional trigonometric parallax has long been used to calculate accurate distances to nearby stars. The regular nature of the parallactic motion of a star, caused by the Earth's orbit around the Sun, allows this motion to be decoupled from the intrinsic proper motion of the star in the heliocentric rest-frame. Hence the distance to the star can be calculated. However, the maximum baseline generating such parallaxes is obviously limited to 2AU. For a given level of astrometric precision, this imposes a fundamental limit to the observable distance. Indeed, the accuracy of parallaxes reported by the Hipparcos mission data (van Leeuwen, 2007) falls to 20–30 percent at best for distances $\sim 300\,\mathrm{pc}$ and only then for the brightest stars. Upcoming astrometric projects such as Pan-STARRS (Kaiser et al., 2002), LSST (Tyson, 2002) and the Gaia mission (Perryman et al., 2001) will achieve similar uncertainty for Sun-like stars as distant as a few kpc, and at fainter magnitudes than was possible with Hipparcos. This extended range will encompass less than 1 percent of the total number of such stars in our Galaxy.

It is clear that it will *not* soon be possible to calculate distances to many of the stars in our Galaxy with conventional trigonometric parallaxes. Alternative means to compute distances to stars are therefore required. Photometry can be used to estimate the absolute magnitude of a star which, when combined with its observed magnitude, allows its distance to be computed. Unfortunately, all attempts to calculate such photometric distances are hindered by the same problems: obscuration by intervening matter alters both observed magnitude (Vergely et al., 1998) and colour (Schlegel et al., 1998; Drimmel & Spergel, 2001), and it is difficult to model appropriate corrections without a reference distance scale. The effects

of chemical composition and age further complicate matters (Jurić et al., 2008). It is therefore difficult to compute photometric distances with an accuracy much better than 20 percent, even for nearby stars, and distances to faint stars are less accurate still (Jurić et al., 2008).

Consequently, there exists a demand for long-range distance measuring tools that can complement the existing tools by overcoming some of their limitations. The work of this chapter, much of which was published in articles by Eyre & Binney (2009a) and Eyre (2010), describes such a tool, which measures distances by utilizing the parallactic motion between a distant star and the Sun, rather than that between a star and the Earth as in conventional parallax.

In the general case, it is not possible to obtain distances this way, because the parallactic motion of the star and its intrinsic proper motion are inextricably mixed up. However, in the special case where the star can be associated with a stellar stream, its rest-frame trajectory can be predicted from the locations of the other associated stars. Using this trajectory, the proper motion in the Galactic rest-frame can indeed be decoupled from the reflex motion of the Sun, and the component of its motion due to parallax can be computed.

In the work of the preceding chapter it was discovered that by measuring the proper-motion vector of a star belonging to a stream, it is possible to directly recover the distance to that star. This discovery is an embodiment of the Galactic parallax effect: since we assume that the peculiar motion of the star is tangent to its stream, *any* proper motion observed perpendicular to that stream can only be due to the reflex motion of the Sun. If the motion of the Sun in a common rest-frame is already known through alternate means, then a parallax can be computed from this perpendicular motion. A diagram illustrating the effect is shown in Fig. 4.1.

Such Galactic parallaxes have the same geometrical basis as conventional trigonometric parallaxes, and as such are free from errors induced by obscuration and reddening. However, the range of Galactic parallax significantly exceeds that of conventional parallax. This is because the Sun orbits about the Galactic centre much faster than the Earth orbits the Sun, and because, unlike with conventional parallax, the Galactic parallax effect is cumulative with continued observation. In realistic cases, for a typically oriented stream, we can expect the Galactic parallax to be observable at nearly 40 times the distance of the equivalent trigonometric parallax, based on 3 years of observations. For increased range, one simply observes over a longer baseline.

This large range means that Galactic parallax might prove a powerful tool to complement conventional parallaxes, and validate other distance measuring tools. It is exciting to note that the capabilities of astrometric projects such as LSST, which will observe the conventional parallax of a G star at a distance of $\sim 1\,\mathrm{kpc}$ with 20 percent uncertainty, will put much of the Galaxy in range of Galactic parallax calculations with similar accuracy.

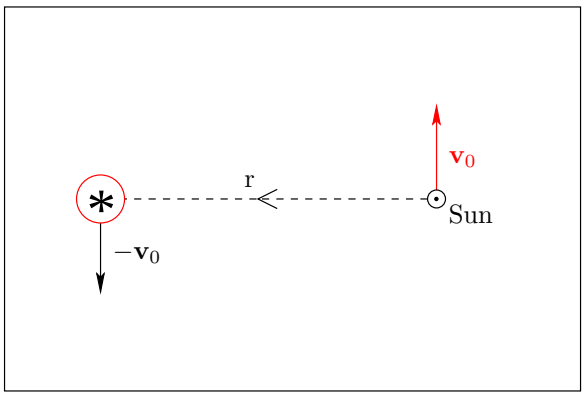

Figure 4.1: An illustration of the Galactic parallax effect. The Sun is at some distance r from a star in a stream. The stream track is oriented perpendicularly to the plane of the diagram. In the rest-frame of the Galaxy, the Sun has a velocity $\mathbf{v}_0$ (marked in red), while the velocity of the stream star points vertically out of the page, consistent with it travelling along the stream. In the heliocentric rest-frame, however, the stream star acquires a velocity component $-\mathbf{v}_0$ (marked in black) in the plane of the diagram, due to the reflex motion of the Sun. Observationally, this reflex motion causes the stream star to acquire a component of proper motion, equal to $v_0/r \equiv \Pi v_0$ projected onto the plane of the sky, that is perpendicular to the stream track. This is the Galactic parallax effect, and allows Π to be calculated if $\mathbf{v}_0$ and the proper motion are known.

The main restriction on the use of Galactic parallax is the requirement for stars to be part of a stream. However, the continuing discovery of significant numbers of streams (Odenkirchen et al., 2003; Majewski et al., 2003; Yanny et al., 2003; Belokurov et al., 2007; Grillmair, 2006a; Grillmair & Dionatos, 2006; Grillmair & Johnson, 2006; Grillmair, 2009; Newberg et al., 2009) using optical surveys implies that they are a staple feature of the Galactic environment, rather than a rarity. The deep surveys of Pan-STARRS and LSST are likely to find yet more, increasing the number of applications for Galactic parallax.

The remainder of this chapter explores the viability of using Galactic parallax to estimate distances and demonstrates its practicality by applying it to data for the GD-1 stream (Grillmair & Dionatos, 2006) published by Koposov et al. (2010, herein K10); we choose to work with the latter over the earlier analysis of the same stream by Willett et al. (2009) on account of the significantly smaller proper-motion uncertainties ($1\,\mathrm{mas\,yr^{-1}}$ vs $4\,\mathrm{mas\,yr^{-1}}$) cited in the later work. Throughout the chapter, the Solar motion is assumed to be

$$(U, V, W) = (10.0, 252, 7.1) \pm (0.3, 11, 0.34)\,\mathrm{km\,s^{-1}},$$

consistent with Aumer & Binney (2009), Reid & Brunthaler (2004) and Gillessen et al. (2009).

The chapter is arranged as follows. §4.2 outlines the calculation of Galactic parallax from observational data. §4.3 examines how uncertainties affect Galactic parallax calculations, and §4.4 discusses the practicality of using Galactic parallax as a distance measuring tool in light of those findings. §4.5

demonstrates the use of Galactic parallax on pseudo-data, and shows that our uncertainty estimates are correct, and that the technique is practical. §4.6 uses Galactic parallax to find distances to the tidal stream GD-1. §4.7 presents our concluding remarks.

4.2 Galactic parallax

Suppose that a star is part of a stellar stream, and has a location relative to the Sun described by $(\mathbf{x} - \mathbf{x}_0) = r\hat{\mathbf{r}}$, where r is the distance to the star, and $\mathbf{x}_0$ is the position of the Sun. In the plane of the sky, let the tangent to the trajectory of the stream, near the star, be indicated by the vector $\hat{\mathbf{p}}$. Assume the velocity of the Sun, $\mathbf{v}_0$, in the Galactic rest-frame (grf) is known, and for the time being assume that streams perfectly delineate orbits. Chapter 3 showed that if the measured proper motion of the star is $\mu\hat{\mathbf{t}}$, then (equation 3.7)

$$\dot{u}\hat{\mathbf{p}} = \mu\hat{\mathbf{t}} + \frac{\mathbf{v}_s}{r} = \mu\hat{\mathbf{t}} + \Pi\mathbf{v}_s, \tag{4.1}$$

where $\Pi \equiv 1/r$ is the Galactic parallax, $\dot{u}$ is the proper motion as would be seen from the grf, and $\mathbf{v}_s$ is the Sun's velocity projected onto the plane of the sky. We note that $\dot{u} = v_t/r$, where v_t is that component of the star's grf velocity perpendicular to the line of sight, and that

$$\mathbf{v}_s = (\mathbf{v}_0 - \hat{\mathbf{r}} \cdot \mathbf{v}_0\, \hat{\mathbf{r}}). \tag{4.2}$$

Equation (4.1) is a vector expression and can be solved simultaneously for both $\dot{u}$ and Π provided that $\hat{\mathbf{p}}$, $\hat{\mathbf{t}}$ and $\mathbf{v}_s$ are not parallel. The stream direction $\hat{\mathbf{p}}$ will not typically be known outright, but must be estimated from the positions of stream stars on the sky. We can achieve this by fitting a low order curve through the position data, the tangent of which is then taken to be $\hat{\mathbf{p}}$. The curve must be chosen to reproduce the gross behaviour of the stream, but we must avoid fitting high-frequency noise, because $\hat{\mathbf{p}}$ is a function of the derivative of this curve, which is sensitive to such noise.

4.3 Uncertainty in Galactic parallax calculations

We begin by taking the cross-product of equation (4.1) with $\hat{\mathbf{p}}$ in order to eliminate $\dot{u}$, and we render the resulting equation into an orthogonal on-sky coordinate system, whose components are denoted by (x, y). Re-arranged for Π, we find

$$\Pi = \frac{\mu\,(t_x \sin\alpha - t_y \cos\alpha)}{v_{s,y}\cos\alpha - v_{s,x}\sin\alpha}, \tag{4.3}$$

where the (x, y) suffixes denote the corresponding components of their respective vectors, and where we have defined the angle $\alpha \equiv \arctan(p_y/p_x)$.

The choice of coordinates (x, y) is arbitrary. We are therefore free to choose the coordinate system in which $\alpha = 0$, i.e. that system in which the x-axis points along the stream trajectory, $\hat{\mathbf{p}}$. Equation (4.3) becomes

$$\Pi = -\frac{\mu t_\perp}{v_{s\perp}}, \tag{4.4}$$

where we now identify the y-component of the various vectors as that component perpendicular ($\perp$) to the stream trajectory, and the x-component as that component parallel ($\|$) to the trajectory. Equation (4.4) shows explicitly that the Galactic parallax effect is due to the reflex motion of stream stars perpendicular to the direction of their travel.

Uncertainties in the (x, y) components of the measured quantities $\mu\hat{\mathbf{t}}$ and $\mathbf{v}_s$, and uncertainty in α, can be propagated to Π using equation (4.3). When we set $\alpha = 0$, this equation becomes

$$\frac{\sigma_\Pi^2}{\Pi^2} = \frac{\sigma_\mu^2}{\mu^2 t_\perp^2} + \frac{\sigma_{v_{s\perp}}^2}{v_{s\perp}^2} + \frac{\sigma_\alpha^2}{v_{s\perp}^2}\left(v_{s\|} + \frac{\mu t_\|}{\Pi}\right)^2, \tag{4.5}$$

where we anticipate the uncertainty in $\mu\hat{\mathbf{t}}$ to be isotropic, and so we have set $\sigma_{\mu t_x} = \sigma_{\mu t_y} = \sigma_\mu$.

We assume σ_μ to be known from observations; it may contain any combination of random and systematic error. $\sigma_{v_{s\perp}}$ is calculated directly from the error ellipsoid on $\mathbf{v}_0$, which is assumed known. Any error on $\mathbf{v}_0$ affects all data in exactly the same way. However, the projection of error on $\mathbf{v}_0$ to $\mathbf{v}_s$ varies with position on the sky. Hence, the effect of $\sigma_{\mathbf{v}_s}$ is to produce a systematic error in reported distance that varies along the stream in a problem-specific way.

Uncertainty in α arises from three sources. Firstly, a contribution $\sigma_{\alpha,m}$ arises due to the possible misalignment of the stream direction with the rest-frame proper motions of the stars. The magnitude of this misalignment is dependent upon both the geometry of the stream and the nature of the host potential, as is detailed extensively in Chapter 5 of this thesis. However, the misalignment is for the most part predictable, given an initial estimate of the stream's orbit. The tangent vectors $\hat{\mathbf{p}}$ can therefore be corrected for the most part of this misalignment, and the residual error is expected to be small.

The second contribution arises because the on-sky trajectory $\hat{\mathbf{p}}$ is chosen by fitting a smooth curve through observational fields, so $\hat{\mathbf{p}}$ need not be exactly parallel to the underlying stream. Further, since $\hat{\mathbf{p}}$ depends on the derivative of the fitted curve, it is likely to be much less well constrained for the data points at the ends of the stream than for those near the middle.

We can quantify this effect. At the endpoints, the fitted curve is likely to depart from the stream by

at most $\Delta\psi$, the angular width of the stream on the sky. For a low-order curve, this departure is likely to have been gradual over approximately half the angular stream length, $\Delta\theta$, giving a contribution to σ_α from fitting of

$$\sigma^2_{\alpha,f} = \frac{4\Delta\psi^2}{\Delta\theta^2}. \tag{4.6}$$

The third contribution to σ_α arises as follows. Since the stream has finite width, at any point, the stars within it have a spread of velocities, corresponding to the spread in action of the orbits that make up the stream. If the stars in a stream show a spread in velocity $(\sigma_{v_x}, \sigma_{v_y})$ about a mean velocity $\mathbf{v_t} = r\dot{u}\hat{\mathbf{p}}$, this effect contributes

$$\sigma^2_{\alpha,v} = \frac{1}{v_t^2}\left(\sigma^2_{v_y}\cos^2\alpha + \sigma^2_{v_x}\sin^2\alpha\right), \tag{4.7}$$

to the uncertainty in α for a single star. Again we can choose $\alpha = 0$, such that $\sigma_{v_y} = \sigma_{v\perp}$, the velocity dispersion perpendicular to the stream direction. Equation (4.7) becomes

$$\sigma^2_{\alpha,v} = \frac{\sigma^2_{v\perp}}{v_t^2} = \frac{\sigma^2_{v\perp}}{(v\sin\beta)^2}, \tag{4.8}$$

where we have introduced v, the grf speed of the stream, and β, the angle of the stream to the line of sight. $\sigma_{v\perp}$ has its origin in the random motions of stars that existed within the progenitor object. In fact, if we assume the stream has not spread significantly in width, then the width and the velocity dispersion (Binney & Tremaine, 2008, §8.3.3) are approximately related by

$$\frac{\sigma_{v\perp}}{v} \simeq \frac{w}{R_\mathrm{p}} = \frac{r\Delta\psi}{R_\mathrm{p}}, \tag{4.9}$$

where w is the physical width of the stream, and R_p is the radius of the stream's perigalacticon. This gives

$$\sigma_{\alpha,v} = \frac{r\Delta\psi}{R_\mathrm{p}\sin\beta}. \tag{4.10}$$

If secular spread has made the stream become wider over time, then this relation will over-estimate $\sigma_{\alpha,v}$, since $\sigma_{v\perp}/v$ is roughly constant. $\Delta\psi$ therefore represents an upper bound on the true value of $\sigma_{\alpha,v}$ through this relation. This argument also assumes that the stream was created from its progenitor in a single tidal event. Real streams do not form in this way. However, repeated tidal disruptions can be viewed as a superposition of ever younger streams, created from a progenitor of ever smaller $\sigma_{v\perp}$.

Equation (4.10) holds for each of these individually. Thus, $\Delta\psi$ remains a good upper bound for $\sigma_{\alpha,v}$ through this relation.

In reality, we do not measure the proper motion of individual stream stars, but rather the mean motion of a field of N stars. The contribution to σ_α is from the error on this mean. Putting this together with equation (4.6) gives our final expression for σ_α

$$\sigma_\alpha^2 = \frac{\sigma_{\alpha,v}^2}{N} + \sigma_{\alpha,f}^2 + \sigma_{\alpha,m}^2 = \frac{r^2\Delta\psi^2}{NR_\mathrm{p}^2\sin^2\beta} + \frac{4\Delta\psi}{\Delta\theta^2} + \sigma_{\alpha,m}^2. \tag{4.11}$$

We note that the first term represents a random error, while the second and third terms represent systematic errors that will vary with position down the stream. In general, $\sin\beta$ and R_p are *a priori* unknown. We can infer $\sin\beta$ from radial velocity information, either directly where the measurements exist, or indirectly from Galactic parallax distances. Guessing R_p requires assumptions to be made about the dynamics, but in general we expect the ratio $r/R_\mathrm{p} \simeq 1$ or less.

Explicit evaluation of $\sin\beta$ and R_p are not necessary to evaluate the uncertainty if the contribution from $\sigma_{\alpha,v}$ is overwhelmed by the error from fitting, $\sigma_{\alpha,f}$. We can see this will be the case when the number of observed stars per field

$$N > \left(\frac{r\Delta\theta}{2R_\mathrm{p}\sin\beta}\right)^2. \tag{4.12}$$

We expect this to be true in almost all practical cases.

4.3.1 Uncertainty in tangential velocity calculations

Equation (4.1) can also be used to solve for $\dot{u}$

$$\dot{u} = \frac{\mu(t_y + t_x) + \Pi(v_{s,y} + v_{s,x})}{\cos\alpha + \sin\alpha}, \tag{4.13}$$

which becomes

$$\dot{u} = \mu(t_\parallel + t_\perp) + \Pi(v_{s\parallel} + v_{s\perp}) = \mu t_\parallel + \Pi v_{s\parallel}, \tag{4.14}$$

when we set $\alpha = 0$. Equation (4.13) combined with equation (4.3) can be used to explicitly propagate uncertainties in the measured quantities to $\dot{u}$. When $\alpha = 0$, the uncertainty in $\dot{u}$ is

$$\begin{aligned}\frac{\sigma_{\dot{u}}^2}{\dot{u}^2} &= \frac{\mathbf{v}_s^2\sigma_\mu^2}{\mu^2(t_\perp v_{s\parallel} - t_\parallel v_{s\perp})^2} + \frac{t_\perp^2(v_{s\perp}^2\sigma_{v_{s\parallel}}^2 + v_{s\parallel}^2\sigma_{v_{s\perp}}^2)}{v_{s\perp}^2(t_\perp v_{s\parallel} - t_\parallel v_{s\perp})^2} \\ &\quad - \frac{2t_\perp^2 v_{s\parallel}v_{s\perp}\mathrm{cov}(v_{s\parallel}, v_{s\perp})}{v_{s\perp}^2(t_\perp v_{s\parallel} - t_\parallel v_{s\perp})^2} + \frac{v_{s\parallel}^2\sigma_\alpha^2}{v_{s\perp}^2}. \end{aligned} \tag{4.15}$$

$\sigma_{v_{s\parallel}}$ and $\mathrm{cov}(v_{s\parallel}, v_{s\perp})$ are calculated directly from the error ellipsoid on $\mathbf{v}_0$, which we have assumed known.

4.4 Practicality of Galactic parallax as a distance measuring tool

Using equation (4.4) to eliminate $\mu t_\perp$ from equation (4.5), and taking the dot product of $\hat{\mathbf{p}}$ with equation (4.1) to simplify the last term, we obtain

$$\begin{aligned}\frac{\sigma_\Pi^2}{\Pi^2} &= \frac{1}{v_{s\perp}^2}\Big\{(r\sigma_\mu)^2 + \sigma_{v_{s\perp}}^2 + (r\dot{u})^2\sigma_\alpha^2\Big\} \\ &= \frac{1}{v_{s\perp}^2}\Big\{(r\sigma_\mu)^2 + \sigma_{v_{s\perp}}^2 + v^2\left(\frac{r^2\Delta\psi^2}{R_\mathrm{p}^2 N} + \frac{4\Delta\psi^2}{\Delta\Theta^2} + \frac{\sigma_{\alpha,m}^2}{\sin^2\beta}\right)\Big\}, \end{aligned} \tag{4.16}$$

where we have noted that $r\dot{u} = v_t = v\sin\beta$, and we have related the observed stream length, $\Delta\theta$, to the deprojected length, $\Delta\Theta = \Delta\theta/\sin\beta$. We note that the last term is independent of r, since $r\Delta\psi = w$, $\Delta\psi/\Delta\Theta$ and $\sigma_{\alpha,m}$ are all constant, and that for a stream of given physical dimension, the uncertainty in Π only has dependence upon the angle of the stream β in the misalignment term $\sigma_{\alpha,m}$.

What level of uncertainty does equation (4.16) predict, when realistic measurement errors are introduced? The answer to this is dependent upon both the physical properties of the stream $(R_p, \Delta\psi, \Delta\Theta, v)$ and the geometry of the problem in question $(r, v_{s\perp}, \beta)$.

We progress by assuming 'typical' values for some of these quantities. The average magnitude of $\mathbf{v}_s$ taken over the whole sky is $v_0\,\pi/4$. The average perpendicular component, for a randomly oriented stream, is $2/\pi$ of this value. We therefore assume a typical value for $v_{s\perp}$ of $v_0/2 \sim 120\,\mathrm{km\,s^{-1}}$. We also assume a typical grf velocity equal to the circular velocity, $v = v_c \sim 220\,\mathrm{km\,s^{-1}}$, and a value for the stream angle to the line of sight, $\sin^2\beta \sim 0.5$, consistent with a randomly oriented stream.

McMillan & Binney (2010) recently summarised the current state of knowledge of $\mathbf{v}_0$. The uncertainty quoted is typically ~ 5 percent on each of (U, V, W). Correspondingly, we estimate a typical value for the uncertainty $\sigma_{v_{s\perp}}$ of 5 percent of $v_{s\perp}$, or $6\,\mathrm{km\,s^{-1}}$, which is dominated by the error on V.

The GD-1 stream that we consider below is exceptionally thin and long, with $\Delta\psi \sim 0.1°$ and $\Delta\theta \sim 60°$.

The Orphan stream (Grillmair, 2006a; Belokurov et al., 2007) is of similar length, but about 10 times thicker. We therefore take $\Delta\psi \sim 1^\circ$, $\Delta\Theta \sim 60^\circ$ as typical of the streams to which one would apply this method. If hundreds of stars are observed for each proper motion datum, then equation (4.12) is true for all realistic combinations of (r, R_p), so we can ignore the contribution of $\sigma_{\alpha,v}$ to σ_α. The contribution from fitting error is $\sigma_{\alpha,f} \simeq 1.9^\circ$. The results of Chapter 5 will show that stream misalignment is highly dependent upon the specifics of the problem in question, but we find that in spherical logarithmic potentials, the typical uncertainty due to misalignment $\sigma_{\alpha,m} \ll 1^\circ$. If the misalignment were larger, then the results of Chapter 5 could be used to make a correction to the tangent vectors derived from the stream. In either case, the residual errors $\sigma_{\alpha,m} \ll 1^\circ$. Hence, we will take $\sigma_\alpha \simeq 2^\circ$ as our typical value.

The individual SDSS-USNO proper motions (Munn et al., 2004) used by K10 have a random uncertainty $\sigma_\mu \sim 4\,\mathrm{mas\,yr^{-1}}$. After averaging over hundreds of stars and accounting for a contribution from non-stream stars, K10 report a random uncertainty of $\sigma_\mu \sim 1\,\mathrm{mas\,yr^{-1}}$ on their GD-1 data. For a stream 10 kpc distant, with these proper motions and the typical values mentioned, equation (4.16) reports an uncertainty of $\sigma_\Pi/\Pi \sim 40$ percent. By far the greatest contribution comes from the first term in equation (4.16), hence, the error on proper-motion measurement is dominating our uncertainty.

To obtain an uncertainty of $\sigma_\Pi/\Pi < 20$ percent with Munn et al. (2004) proper-motion measurements, we would need to restrict ourselves to streams less than 5 kpc distant. 20 percent error is also possible at 10 kpc given optimum problem geometry. This is clearly competitive with the ~ 20 pc at which one could observe a standard trigonometric parallax, with similar accuracy, using astrometry of this quality. However, previous work (Willett et al., 2009, K10) shows that SDSS photometry combined with population models produce distance estimates accurate to ~ 10 percent for stars in streams at 8 kpc. The accuracy of Galactic parallax is therefore not likely to be as good as that of photometric distances for distant streams, using data this poor, unless the problem geometry is favourable.

Proper-motion data from the Pan-STARRS telescope is expected to be accurate to $\sim 1\,\mathrm{mas\,yr^{-1}}$ for Sun-like stars at 10 kpc (Magnier et al., 2008). K10 reduce raw data with accuracy $\sim 4\,\mathrm{mas\,yr^{-1}}$ to processed data accurate to $\sim 1\,\mathrm{mas\,yr^{-1}}$, even though the expected proper motion of the stars is of the same size as the errors. It is not unreasonable to expect a similar analysis applied to Pan-STARRS raw data, where the relative error would be much less than unity, to yield processed data accurate to $\sim 0.2\,\mathrm{mas\,yr^{-1}}$. In truth, the ability of Pan-STARRS to detect very faint stars will increase the number of stars identifiable with a stream, and thus reduce the uncertainty in the mean proper motion further than this, but we use $0.2\,\mathrm{mas\,yr^{-1}}$ as a conservative estimate.

The same 10 kpc distant stream would have a parallax error of $\sigma_\Pi/\Pi \simeq 11$ percent with data this accurate. An error of less than 20 percent is possible for a typical stream less than ~ 23 kpc distant, and for a stream with favourable geometry less than 50 kpc distant. Jurić et al. (2008) report that

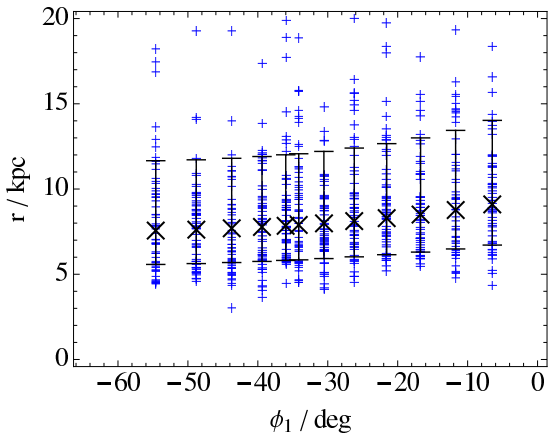

Figure 4.2: Crosses: Galactic parallax distances computed from the pseudo-data, with no extraneous errors. Error bars: the random scatter expected in Galactic parallax distances, with measurement errors as mentioned in the text. Plus signs: Galactic parallax distances computed from 60 Monte Carlo realisations of each pseudo-datum convolved with the measurement errors. The analytic uncertainty estimate and the Monte Carlo realisations are in good agreement.

SDSS photometric distances for dwarf stars have ~ 40 percent error at 20 kpc. Thus, the accuracy of Galactic parallax derived from Pan-STARRS data should be at least comparable to distance estimates from photometric methods, even in the typical case.

Future projects such as LSST and Gaia will each obtain proper motions accurate to $\sim 0.2\,\mathrm{mas\,yr^{-1}}$ for Sun-like stars 10 kpc distant (Ivezić et al., 2008; Perryman et al., 2001). These data would allow a distance estimate for our typical stream accurate to 8 percent, and a stream with favourable geometry accurate to 4 percent. Error in the proper motion no longer dominates the uncertainty in these calculations. We might expect such accurate astrometric surveys to reduce the uncertainty in the Solar motion; in this case, the error in Galactic parallax would be lower still.

Gaia will not observe Sun-like stars beyond 10 kpc, but LSST will, with accuracy of $0.4\,\mathrm{mas\,yr^{-1}}$ for dwarf stars 30 kpc distant (Ivezić et al., 2008). The accuracy of the parallax to our typical stream at this distance would be about 14 percent with these data, and 6 percent is achievable with optimum geometry. A typical stream could be measured to 20 percent accuracy out to 40 kpc, and a stream with favourable geometry out to 54 kpc; this range approaches the limit of LSST's capability for detection of dwarf stars. Such data will put the Orphan stream, which is about $20-30$ kpc distant (Grillmair, 2006a; Belokurov et al., 2007; Sales et al., 2008), in range of accurate trigonometric distance estimation. For comparison, photometric distances from SDSS data are hardly more accurate than 50 percent for this stream (Belokurov et al., 2007).

4.5 Tests of Galactic parallax

To test the method, pseudo-data were prepared from an orbit fitted to data for the GD-1 stream by K10. The orbit is described by the initial conditions

$$\mathbf{x} = (-3.41, 13.00, 9.58)\,\mathrm{kpc}, \qquad \mathbf{v} = (-200.4, -162.6, 13.9)\,\mathrm{km\,s^{-1}},$$

where the x-axis points towards the Galactic centre, and the y-axis points in the direction of Galactic rotation. The orbit was integrated in the logarithmic potential

$$\Phi(x, y, z) = \frac{v_c^2}{2} \log\left(x^2 + y^2 + \left(\frac{z}{q}\right)^2\right), \tag{4.17}$$

where $v_c = 220\,\mathrm{km\,s^{-1}}$ and $q = 0.9$. The resulting trajectory was projected onto the sky, assuming a Solar radius $R_0 = 8.5\,\mathrm{kpc}$. The implicit assumption is made that the GD-1 stream is closely approximated by an orbit: the work of Chapter 5 will show that this assumption is a fair one. Several points were sampled from the on-sky track, and each was taken to be a separate datum in the pseudo-data set. The proper motion for each datum was computed by projecting the difference between its grf motion and the Solar motion onto the sky.

The pseudo-data were transformed into the rotated coordinate system used by K10 to facilitate comparison with their data; the transformation rule is given in the appendix to K10. The stream is very flat in this coordinate system, so the dependence of ϕ_2 on ϕ_1 is relatively weak. This helps to increase the quality of the fitted curve and minimises the corresponding error in $\sigma_{\alpha,f}$.

To simulate the observed scatter in the real positional data, the pseudo-data were each scattered in the ϕ_2 coordinate according to a randomly-sampled Gaussian distribution with a dispersion $\sigma_{\phi_2} = 0.1^\circ$. The resulting positional pseudo-data are plotted in Fig. 4.3, along with the orbit from which they were derived (full curve). A cubic polynomial representing $\phi_2(\phi_1)$ was least-squares fitted to the pseudo-data, the tangent of which was used to estimate $\hat{\mathbf{p}}$. In the case of the pseudo-data, uniform weights were applied to each datum for the fitting processes. The resulting curve is also shown in Fig. 4.3 (dotted curve).

When the correct orbit is used to calculate $\hat{\mathbf{p}}$, and precise values for the measured proper motion $\mu\hat{\mathbf{t}}$ and Solar reflex motion $\mathbf{v}_s$ are used, the distance is recovered perfectly from equation (4.4). Fig. 4.4 compares the recovered distance when $\hat{\mathbf{p}}$ is estimated using the polynomial fit to the pseudo-data, but still using accurate values for $\mu\hat{\mathbf{t}}$ and $\mathbf{v}_s$. Our pseudo-data stream is $\Delta\psi \simeq 0.1^\circ$ wide and $\Delta\theta \simeq 60^\circ$ long. Equation (4.6) therefore estimates $\sigma_{\alpha,f} \simeq 0.38^\circ$. The recovered distances in Fig. 4.4 are in error by only ~ 2 percent across most of the range, which is the approximate uncertainty predicted by equation (4.5)

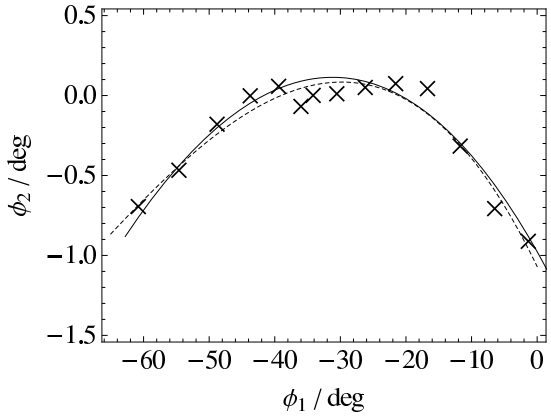

Figure 4.3: Full line: the orbit for the GD-1 stream, taken from K10. Crosses: pseudo-data derived from that orbit, but randomly scattered in ϕ_2 according to a Gaussian distribution with a dispersion of $\sigma_{\phi_2} = 0.1°$. Dotted line: a cubic polynomial fitted to the pseudo-data, used to estimate stream direction.

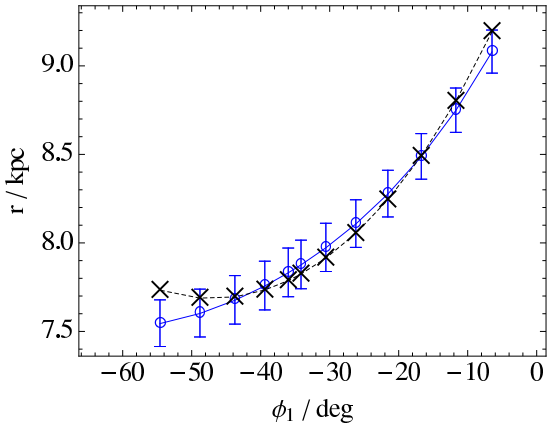

Figure 4.4: Dotted line: the orbit of the GD-1 stream, taken from K10. Crosses: the true distance of each pseudo-datum. Circles: Galactic parallax distances computed from the pseudo-data. The error bars represent the distance error expected from the polynomial fitting procedure. No extraneous error was added. The error bars are shown to be a good estimate of likely error from the fitting procedure, and the agreement of the distances overall is excellent.

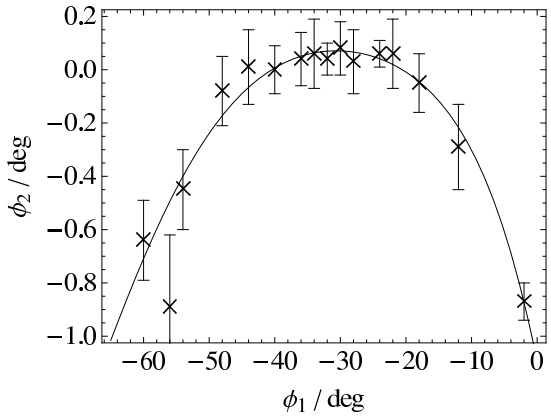

Figure 4.5: Crosses: on-sky position data for the GD-1 stream, as published in K10. The error bars represent the quoted uncertainties. Full line: linear least-squares fit of a cubic polynomial, $\phi_2(\phi_1)$, to these data; the inverse-square of the uncertainties was used to weight the fit.

for this value of $\sigma_{\alpha,f}$. Thus, the estimation of $\hat{\mathbf{p}}$ from the observed stream is good, and contributes little error to the distance calculations.

The K10 observational data for the GD-1 stream, discussed below, have a similar uncertainty $\sigma_\alpha \sim 0.38^\circ$ due entirely to the fitting process, and proper-motion uncertainties $\sigma_\mu \sim 1\,\mathrm{mas\,yr}^{-1}$. Fig. 4.2 shows the recovered distances from Fig. 4.4 with error bars for the expected uncertainty in recovered distance, given these measurement uncertainties and the uncertainty in $\mathbf{v}_0$ quoted in §4. Also plotted for each datum are the distances recovered from 60 Monte Carlo realisations of the pseudo-data input values, convolved with the errors given above.

Equation (4.5) is found to be a good estimator for the uncertainty, with approximately 80 percent of the Monte Carlo realisations falling within the error bars. The error in parallax for the K10 data is thus predicted to be about 50 percent, of which the greatest contribution comes from the uncertainty in proper motion.

4.6 Distance to the GD-1 stream

Fig. 4.5 shows the on-sky position data for the GD-1 stream, as published in K10. Also shown in Fig. 4.5 is a linear least-squares fit of a cubic polynomial to these data, used to estimate $\hat{\mathbf{p}}$. The weights for the fit were the inverse-square uncertainties for each position field, as given by K10.

K10 provide measured proper-motion data for five fields of stars, spanning the range $\phi_1 \sim (-55, -15)^\circ$, along with uncertainties for these measurements. The uncertainty in $\mathbf{v}_s$ is computed for each individual field from the uncertainty in $\mathbf{v}_0$ given in §4. Uncertainty in the stream direction is $\sigma_\alpha \sim 0.38^\circ$, which is entirely contributed by the curve fit to the stream; since hundreds of stars contributed to the calculation of the proper motions, the contribution from the first term in equation (4.11) is negligible. We have

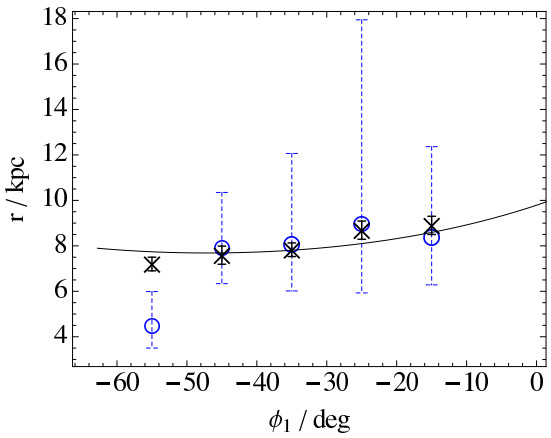

Figure 4.6: Circles: Galactic parallax distances for the GD-1 data presented in K10. Dotted error bars: the uncertainty estimated by equation (4.5), given the K10 measurement uncertainties. Crosses: the photometric distances reported in K10, along with their error bars. Full line: the orbit for GD-1 taken from K10. With the exception of the datum at $\phi_1 \sim -55$ deg, the Galactic parallax distances are in excellent agreement with the photometric distances from K10. The dotted error bars appear to seriously over-estimate the true error in the distance estimates.

also assumed that the contribution to σ_α from stream misalignment is negligible. This assumption is validated by the results of Chapter 5, which predicts that the projection of GD-1 should be perfectly aligned with the orbits of its stars.

Fig. 4.6 shows the Galactic parallax distances for each of these data, along with the K10 photometric distances. The dotted error bars represent the expected error in distance for the uncertainties given. The small solid error bars are the uncertainties reported by K10 for their photometric distances. The K10 orbit used to compute the earlier pseudo-data is plotted for comparison.

With the exception of the datum at $\phi_1 \sim -55°$, the parallax distances and the K10 distances are in remarkable agreement. However, the dotted error bars vastly overestimate the true error in the results. If we ignore the datum at $\phi_1 \sim -55°$, the scatter in the distance, $\sigma_r \sim 1\,\mathrm{kpc}$, is similar to that of the photometric distances, and consistent with a true random error of $\sigma_\mu \sim 0.3\,\mathrm{mas\,yr^{-1}}$, and negligible systematic offset. We cannot explain this discrepancy, except by suggesting that the K10 proper-motion measurements are more accurate than the published uncertainties suggest. This is corroborated by the top-right panel of Fig. 13 from K10 in which the μ_{ϕ_2} data, with the exception of the datum at $\phi_1 \sim -55°$, show remarkably little scatter within their error bars.

Fig. 4.7 shows the Galactic rest-frame proper motions, $\dot{u}$, calculated from equation (4.14) along with their error bars, from equation (4.15). In the background are plotted the data from Fig. 9 of K10, which show the density of stars with a given grf proper motion in the sample of stars chosen to be candidate members of the stream, and after subtraction of a background field. The K10 grf proper motions have been calculated by correcting measured proper motion for the solar reflex motion, using an assumed

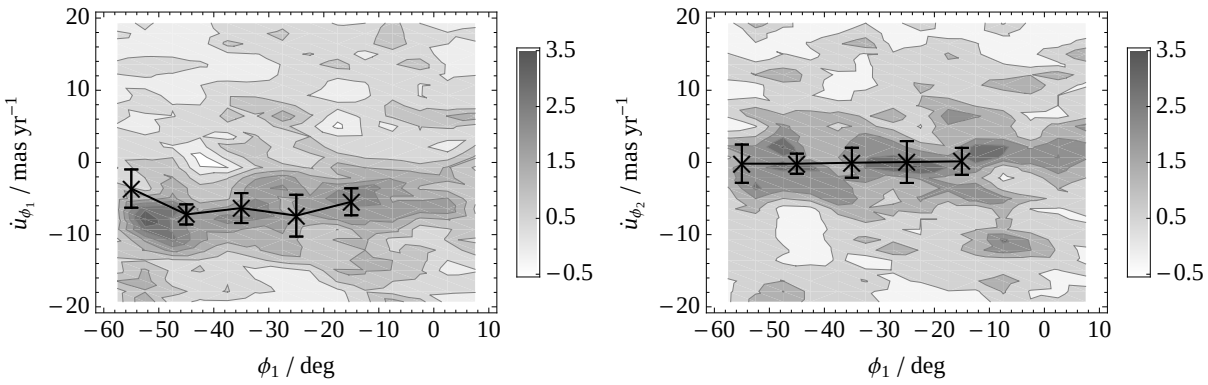

Figure 4.7: Full lines: Galactic rest-frame proper motion, $\dot{u}$, calculated from the K10 data using equation (4.14). The (left, right) panels show the (ϕ_1, ϕ_2) components respectively. Plotted in the background are the observational data from Fig. 9 in K10; the greyscale shows the number of stream stars, per bin, with the given motion. The data are broadly consistent, except for the datum at $\phi_1 \sim -55$ deg in the left panel.

distance of 8 kpc (Koposov, private communication); this assumption will cause a systematic error in the K10 proper motions, of order the distance error, which changes with position down the stream. The apparently large width of the stream in this plot is due to uncertainty in the underlying Munn et al. (2004) proper-motion data.

The stream is clearly visible in this plot as the region of high density spanning $\phi_1 \sim (0, -60)^\circ$ with $\dot{u}_{\phi_2} \simeq 0\,\mathrm{mas\,yr}^{-1}$ and $\dot{u}_{\phi_1}$ falling slowly between $(-6, -10)\,\mathrm{mas\,yr}^{-1}$. Despite the expected systematic error, the estimates of $\dot{u}$ from the parallax calculation are consistent with these data, with the exception of the same datum at $\phi_1 \sim -55^\circ$ that also reports an anomalous distance.

We explain this suspect datum as follows. From inspection of the top-right panel of Fig. 13 from K10, it is apparent that the μ_{ϕ_2} measurement for this datum is not in keeping with the trend. Conversely, the corresponding μ_{ϕ_1} measurement is not obviously in error. If the magnitude of μ_{ϕ_2} for this datum has been over-estimated by the K10 analysis, then equation (4.4) will over-estimate the parallax, and hence under-report the distance. Fig. 4.6 indicates that the distance for this datum is indeed under-reported.

The effect of such an error in μ_{ϕ_2} on the grf proper-motion, $\dot{u}$, can be understood by considering equation (4.14). If Π is over-estimated, $\dot{u}$ will be either over-estimated or under-estimated, depending on the relative sign of the two terms. In the case of GD-1, $\mu t_\parallel$ and $v_{s\parallel}$ have opposite signs, so an over-estimated Π will result in an under-estimated $\dot{u}$. This too corresponds with the behaviour of the suspect datum in Fig. 4.7.

It is unknown why this particular datum should be significantly in error while the other data are not. There are no obvious structures in the lower panel of Fig. 4.7 which might cause the fitting algorithm in K10 to mistakenly return an incorrect value for μ_{ϕ_2}. Nonetheless, if the scatter in the other data

are accepted as indicative of their true statistical error, it is clear that the datum at $\phi_1 \sim -55^\circ$ cannot represent the proper motions of GD-1 stars at that location. We therefore predict that an appropriate re-analysis of the proper-motion data, taking care to ensure that a signal from GD-1 stream stars is properly detected, will return a revised proper-motion of $\mu_{\phi_2} \sim -3\,\mathrm{mas\,yr^{-1}}$.

In summary, it seems that Galactic parallax measurements confirm the K10 photometric analysis, and predict that the stream is approximately $(8 \pm 1)\,\mathrm{kpc}$ distant, where the uncertainty denotes the scatter in the results. Since Galactic parallax and photometric estimates are fundamentally independent, it seems unlikely that systematic errors in either would conspire to produce the same shift in distance; this suggests that no systematic error is present.

We also calculate a grf proper motion for the stream of $\dot{u}_{\phi_1} = (-7 \pm 2)\,\mathrm{mas\,yr^{-1}}$, corresponding to a grf tangential velocity of $(265 \pm 80)\,\mathrm{km\,s^{-1}}$ in a direction $(\dot{u}_l \cos b, \dot{u}_b) \simeq (0.8, -0.6)$. This implies that the stream is on a retrograde orbit, inclined to the Galactic plane by $\sim 37^\circ$, which is in accordance with previous results (Willett et al., 2009; Koposov et al., 2010).

The galactocentric radius of $\sim 14.5\,\mathrm{kpc}$ does not seem to be changing rapidly along the stream's length, which subtends $\sim 12^\circ$ when viewed from the Galactic centre. This implies that the observed stream is at an apsis. The grf velocity of the stream is faster than the circular velocity, $v_c \sim 220\,\mathrm{km\,s^{-1}}$. This implies that the stream is at pericentre, although the large uncertainty prevents a firm conclusion from being drawn. We note that the radial velocity data in K10 would also imply that the stream is observed at pericentre.

4.7 Conclusions

In this chapter, we have demonstrated the practical application of a technique for computing distances using Galactic parallax. This technique utilises the predictable trajectories of stars in a stream to identify the contribution of the reflex motion of the Sun to the observed proper motion. The parallax and the Galactic rest-frame proper motion follow from this.

Galactic parallax is a geometric phenomenon, and the distances obtained from it are in every way as fundamental as those obtained using conventional trigonometric parallax. The only assumption made is knowledge of the Galactic rest-frame velocity of the Sun. It is also a requirement that the observed stars are part of a stream. Recent evidence (Odenkirchen et al., 2003; Majewski et al., 2003; Yanny et al., 2003; Belokurov et al., 2007; Grillmair, 2006a; Grillmair & Dionatos, 2006; Grillmair & Johnson, 2006; Grillmair, 2009; Newberg et al., 2009) indicates that tidal streams are a common constituent of the Galactic halo, and so this technique should have widespread application.

The key hurdle to the widespread application of this technique is the lack of high-accuracy proper-

motion data for distant stream stars. Given the advent of next-generation astrometric projects such as Pan-STARRS (Kaiser et al., 2002), LSST (Tyson, 2002) and Gaia (Perryman et al., 2001), proper-motion catalogues with billions of entries of the required precision will become available in the next few years. The technique should then become the standard method of determining the distance to remote stream stars.

4.7.1 Summary

We have presented a method for calculating the Galactic parallax of tidal streams. We first determine the on-sky direction of the stream by fitting it with a curve. We then combine the tangents of this curve with measured proper-motion data to estimate the parallax of the stream.

We have derived an expression for the uncertainty in Galactic parallax calculations. We include contributions from measurement errors in proper motion and the Solar motion, error in the estimation of stellar trajectories from the stream direction, and algorithmic error in the estimation of stream direction itself.

The uncertainty for calculations involving a particular stream depends upon the size, location and orientation of the stream, as well as upon measurement errors. We estimate that using individual proper motions accurate to $4\,\mathrm{mas\,yr^{-1}}$, available now in published surveys (Munn et al., 2004), the parallax of a 10 kpc distant stream with typical geometry can be measured with an uncertainty of 40 percent. The parallax of a stream with optimum geometry could be measured with approximately half this uncertainty.

Proper-motion data from the forthcoming Pan-STARRS PS-1 survey (Kaiser et al., 2002; Magnier et al., 2008) will yield the distance to a typical 10 kpc distant stream with 11 percent accuracy, or the distance to a stream at 23 kpc with 20 percent accuracy; with favourable geometry this accuracy could be achieved for a stream as distant as 50 kpc. With data of this quality, the uncertainty in distances from Galactic parallaxes will be considerably lower than those of photometric distances to remote streams.

The LSST (Tyson, 2002; Ivezić et al., 2008) and the Gaia mission (Perryman et al., 2001) will produce proper-motion data that are more accurate still. Such data would allow the distance to stars in a 10 kpc typical stream to be computed to an accuracy of order of 8 percent, where the limitation is now imposed by uncertainty in the solar motion and in the stream trajectory. It is likely that LSST and Gaia data will allow the uncertainty in the solar motion to be significantly reduced, so in reality much better precision can be expected at this distance. For streams 30 kpc distant, LSST proper motions will allow distance estimates as accurate as 14 percent to be made in the typical case, and 6 percent with optimum geometry. Thus, the high-quality astrometric data that is expected to be available in the next decade will allow parallax estimates for very distant streams to be made with unparalleled accuracy.

To test the method presented, we have created pseudo-data simulating the GD-1 stream (Grillmair & Dionatos, 2006). When the method is provided with error-free pseudo-data, the correct parallax is computed perfectly. When errors are introduced into the pseudo-data, the reported parallax degrades in line with the uncertainty estimates.

We applied the method to the astrometric data for the GD-1 stream in Koposov et al. (2010). With the exception of a single datum, the Galactic parallax is remarkably consistent with the photometric distances quoted by Koposov et al. (2010). Indeed, the uncertainty in the measured proper motions quoted by Koposov et al. (2010) should produce significant error in the Galactic parallax. However, the scatter in the results is consistent with a random error of only $\sim 0.3\,\mathrm{mas\,yr^{-1}}$, and if the photometric distances of K10 are believed, no systematic offset. This is at odds with the typical uncertainty in the proper motion of $\sim 1\,\mathrm{mas\,yr^{-1}}$ reported by Koposov et al. (2010). We cannot explain this discrepancy, other than to suggest that the Koposov et al. (2010) method for estimating error in the proper motions is producing significant over-estimates.

The Galactic rest-frame proper motions predicted for the stream are also consistent with observational data from Koposov et al. (2010), with the exception of the same datum that also reports an inconsistent distance. We conclude that the proper-motion associated with this datum is erroneous, and we predict that reanalysis of the stream stars near this datum will reveal a reduced proper-motion measurement of $\mu_{\phi_2} \sim 3\,\mathrm{mas\,yr^{-1}}$.

Photometry and Galactic parallax produce fundamentally independent estimates of distance. The quality of the corroboration of the Koposov et al. (2010) photometric distance estimates for GD-1 by the Galactic parallax estimates presented here therefore lends weight to the conclusion that the predicted distance, in both cases, is correct. On this basis, we conclude that the GD-1 stream is about $(8 \pm 1)\,\mathrm{kpc}$ distant from the Sun, on a retrograde orbit that is inclined 37° to the Galactic plane with a rest-frame velocity of $(265 \pm 75)\,\mathrm{km\,s^{-1}}$. We also conclude that the visible portion of the stream is probably at pericentre.

4.7.2 Future directions

The prospect of being able to map trigonometric distances in the Galaxy to high accuracy at a range of tens of kiloparsecs is indeed exciting. The distances generated using this method, although limited to stars in streams, could be used to calibrate other distance measuring tools, such as photometry, that would be more widely applicable. The technique is immediately applicable to any stream for which proper-motion data are currently available, although we anticipate limited accuracy until better proper-motion data are available.

Given enough parallax data points along a given stream, an orbit can be constructed by connecting those points. This orbit is predicted independently of any assumption about the Galactic potential, which it must strongly constrain. Constraints on the Galactic potential impose constraints on theories of galaxy formation and cosmology. It would seem that the combination of dynamics and Galaxy-scale precision astrometry, such as provided by this method, could well have profound implications for astrophysics in the future.

At present, however, it is not obvious how to combine all sources of astrometric and dynamical information, to produce the tightest constraints on the potential. A significant theoretical effort is therefore required to explore methods for combining this information, in anticipation of the arrival of higher quality astrometric data in the next few years.

Chapter 5

The mechanics of streams

5.1 Introduction

Many of the techniques used to harness the diagnostic power of tidal streams rely upon the assumption that such streams delineate orbits precisely (Binney, 2008; Eyre & Binney, 2009a; Odenkirchen et al., 2009; Willett et al., 2009; Koposov et al., 2010; Eyre, 2010).

Although both the author (Chapter 2; Eyre & Binney, 2009b) and others (Dehnen et al., 2004; Choi et al., 2007; Montuori et al., 2007) have shown evidence that this is not necessarily the case, there has not to date been an exposition of how streams form that fundamentally addresses this issue. In particular, it is necessary to understand under what circumstances tidal streams delineate orbits, by what measure they are in error when they do not, and what can be done to correct this error.

The few studies that have been made have either focused on N-body simulations (e.g. Choi et al., 2007; Montuori et al., 2007), the confusion of which makes the predominant physics hard to isolate, or have attempted to describe the problem in terms of conventional phase-space coordinates and classical integrals (e.g. Dehnen et al., 2004; Choi et al., 2007), which makes the problem intractably hard. In particular, the work of Choi et al. (2007) made some progress towards understanding the dynamical structure of clusters at the point of disruption, and they provide a qualitative picture of the evolution of tidal tails, understood in terms of classical integrals. However, they are unable to make predictions for stream tracks on the basis of this picture alone, and they are ultimately forced to rely on N-body simulation. In the work of this chapter, we approach the problem using action-angle variables (Binney & Tremaine, 2008, §3.5), in which the physics of stream formation turns out to have a natural and simple expression.

We confine our investigation to the formation of long, cold streams, such as may form from tidally

disrupted globular clusters. We do so for two reasons. Firstly, a low mass for the progenitor cluster simplifies the understanding of its orbital mechanics, because of the lack of tidal friction and other feedback effects in their interaction with the host galaxy. Secondly, thin, long streams provide the strongest constraints upon the Galactic potential (Chapter 2; Binney, 2008; Eyre & Binney, 2009b), because any orbit delineated by them can be observationally identified with less ambiguity. It is therefore long, cold streams that are of primary interest for use in probing the potential.

We study the mechanics of stream formation immediately following the tidal disruption of a progenitor cluster. In most of the work that follows, the assumption is made that stream stars feel only the potential of the Galaxy; i.e. the stream stars do not self-gravitate. This assumption is generally a fair one: the stars in streams are generally spaced too widely for their self-gravity to be of consequence (Dehnen et al., 2004). Indeed, we will demonstrate in §5.6.2 below that self-gravity becomes negligible shortly after stars are stripped from the cluster.

The chapter is arranged as follows: The remainder of this introduction discusses the action-angle variables in which we perform our analysis. §5.2 discusses the basic mechanics of stream formation and propagation. §5.3 discusses the detail of stream formation in spherical systems, and explores some examples. §5.4 describes the principles and pitfalls of mapping streams from action-angle space to real-space. §5.5 describes the consequences of optimizing potential parameters by assuming that streams follow orbits. §5.6 examines the action-space distribution of disrupted clusters using N-body simulation. §5.7 discusses stream formation in flattened systems, using oblate axisymmetric Stäckel potentials as an example. Finally, §5.8 presents our concluding remarks.

5.1.1 Action-angle variables

We will approach the analysis of stream formation and propagation by describing the problem using action-angle variables. The usefulness and theoretical basis of action-angle variables is extensively discussed in §3.5 of Binney & Tremaine (2008). Here, we merely note that action-angle variables are a set of canonical coordinates, like conventional phase-space coordinates, that can be used to describe systems in Hamiltonian mechanics.

The coordinates are special because the canonical momenta, called actions, $\mathbf{J}$, are integrals of the motion, and are thus constant with the passage of time. The equation of motion for $\mathbf{J}$ then reads

$$\dot{J}_i = 0 = -\frac{\partial H}{\partial \theta_i}, \tag{5.1}$$

where we have introduced the angle variables, $\boldsymbol{\theta}$. Equation (5.1) requires that the Hamiltonian, $H(\mathbf{J})$ be

independent of $\boldsymbol{\theta}$, and therefore a function of $\mathbf{J}$ only. The equation of motion for $\boldsymbol{\theta}$ is then,

$$\dot{\theta}_i = \frac{\partial H}{\partial J_i} \equiv \Omega_i(\mathbf{J}), \tag{5.2}$$

where the frequencies $\boldsymbol{\theta}(\mathbf{J})$, being functions of $\mathbf{J}$ only, are constant. The solution to equation (5.2) is very simple

$$\boldsymbol{\theta}(t) = \boldsymbol{\theta}(0) + \boldsymbol{\Omega}\, t. \tag{5.3}$$

Hence, the motion of a system described by action-angle variables is very easy to predict. Moreover, because the actions for the system are constant, any approximation that we may make to $H(\mathbf{J})$, if valid at one time, is automatically valid at all times. This particular feature is of critical importance in simplifying the analysis of the propagation of streams in the work that follows.

§3.5 of Binney & Tremaine (2008) discusses at length the calculation of action-angle coordinates in various systems. Here, we note that standard methods for calculating action-angle coordinates require the Hamilton-Jacobi equation (equation 3.205, Binney & Tremaine, 2008) to separate. This condition is met by all spherically symmetric potentials, but excludes any asymmetric potential that is not of Stäckel form (see §5.7.2 below). Given this condition, the action corresponding to a coordinate q is given by

$$J_q = \frac{1}{2\pi} \oint p_q \, dq, \tag{5.4}$$

where the integral is over the closed path that encloses a single oscillation of the coordinate q along an orbit.

We note that just as spherical systems are naturally described by spherical polar coordinates (r, ϑ, ϕ), the natural actions for such systems are (J_r, L), where L is the angular momentum, and we have without loss of generality confined our motion to the (x, y) plane. We can qualitatively understand the meaning of these actions. An orbit with finite L we understand to be circular at, or to oscillate in epicycles around, some guiding-centre radius r_g. An orbit with $J_r = 0$ always has zero radial momentum, and thus corresponds to a circular orbit. Conversely, an orbit with comparatively large J_r must be very eccentric. Hence, J_r quantifies the radial motion of an orbit, while L quantifies its azimuthal motion. We note for completeness that orbits with $L = 0$ are plunging orbits.

In the sections that follow, we utilize the equations from §3.5.2 of Binney & Tremaine (2008) to transform between the action-angle coordinates $(J_r, L, \theta_r, \theta_\phi)$ and conventional phase-space coordinates, when investigating spherical potentials. The action-angle coordinates used when investigating non-spherical potentials are discussed in §5.7.2 below.

5.2 The formation of streams from tidally stripped clusters

Consider a low-mass cluster, just past the point of disruption, such that the stripped stars no longer feel the effects of its gravity. The cluster is on a regular orbit, identified by its actions $\mathbf{J}_0$, in a fixed background potential, which has a Hamiltonian H in terms of the actions $\mathbf{J}$. Suppose further that in the locality of $\mathbf{J}_0$, the Hamiltonian is well described by the Taylor expansion,

$$H(\mathbf{J}) = H_0 + \boldsymbol{\Omega}_0 \cdot \delta\mathbf{J} + \frac{1}{2}\delta\mathbf{J}^T \cdot \mathbf{D} \cdot \delta\mathbf{J}, \tag{5.5}$$

where $\delta\mathbf{J} = \mathbf{J} - \mathbf{J}_0$, and $\mathbf{D}$ is the Hessian of H

$$D_{ij} = \left.\frac{\partial^2 H}{\partial J_i\, \partial J_j}\right|_{\mathbf{J}}, \tag{5.6}$$

and $\boldsymbol{\Omega}_0$ is the frequency of the cluster's orbit

$$\Omega_{0,i} = \left.\frac{\partial H}{\partial J_i}\right|_{\mathbf{J}}. \tag{5.7}$$

The frequency $\boldsymbol{\Omega}$ of a nearby orbit $\mathbf{J}$ is then

$$\boldsymbol{\Omega}(\mathbf{J}) = \boldsymbol{\Omega}_0 + \mathbf{D} \cdot \delta\mathbf{J}. \tag{5.8}$$

If the disrupted cluster has some spread in actions $\Delta\mathbf{J}$ and angles $\Delta\boldsymbol{\theta}_0$, then the spread in angles after some time t is given by (Tremaine, 1999; Binney & Tremaine, 2008, §8.3.1),

$$\Delta\boldsymbol{\theta}(t) = t\Delta\boldsymbol{\Omega} + \Delta\boldsymbol{\theta}_0 \simeq t\Delta\boldsymbol{\Omega}, \tag{5.9}$$

where the near equality is valid when $\Delta\boldsymbol{\Omega}\, t \gg \Delta\boldsymbol{\theta}_0$, and where we have introduced the spread in frequencies, $\Delta\boldsymbol{\Omega}$, which are related to the spread in actions via

$$\Delta\boldsymbol{\Omega} = \mathbf{D} \cdot \Delta\mathbf{J}. \tag{5.10}$$

In the absence of self-gravity, the action-space distribution $\Delta\mathbf{J}$ is frozen for all time. $\mathbf{D}$ and $\boldsymbol{\Omega}$ are functions of $\mathbf{J}$ only, and so are similarly frozen. The secular evolution of a disrupted cluster is therefore to spread out in angle-space, with its eventual shape determined by $\Delta\boldsymbol{\Omega} = \mathbf{D} \cdot \Delta\mathbf{J}$, and its size growing linearly with t.

For a given $\Delta\mathbf{J}$, what does $\Delta\boldsymbol{\Omega}$ look like? Equation (5.9) and equation (5.10) show that the function

of $\mathbf{D}$ is to act as a linear map between a star's position in action-space and its position in angle-space. We note that $\mathbf{D}$ is a Hessian and is therefore symmetric. Associated with $\mathbf{D}$ are three orthogonal directions, $\hat{\mathbf{e}}_n$ $(n = 1, 3)$, corresponding to the eigenvectors of $\mathbf{D}$ if it is evaluated as a matrix. Each of these directions is associated with an eigenvalue, λ_n.

Consider a cluster whose stars are distributed isotropically within a unit sphere about some mean $\mathbf{J}_0$ in action-space. The corresponding angle-space structure, resulting from the mapping of this sphere by $\mathbf{D}$, will be an ellipsoid with semi-axes of length $t\lambda_n$ and direction $\hat{\mathbf{e}}_n$.

If the λ_n are finite and approximately equal, such an isotropic cluster will spread out in angle-space with no preferred direction, with the density initially falling as t^3. Eventually, the cluster will uniformly populate the whole of angle-space; in real-space the cluster will uniformly fill the entire volume occupied by the orbit $\mathbf{J}_0$. For a cluster disrupting in an actual galaxy, the structure would quickly fall below the level of observability.

If one of the λ_n is much smaller than the others, then $\mathbf{D}$ will act to map the cluster into a highly flattened ellipsoid in angle-space. In real-space, the cluster will occupy a 2-dimensional subspace of the orbital volume of $\mathbf{J}_0$. The precise configuration of this subspace is likely to be complex. However, the density will initially fall as t^2, and in a real galaxy, such a structure will rapidly become invisible.

If two of the λ_n are small, then $\mathbf{D}$ acts to map the cluster into a line in angle-space. In this case, the resulting real-space structure will be a filament. The density of this filament will fall linearly with t. In a real galaxy, such a structure may therefore persist with a significant overdensity for some time. It is this case that describes the formation of tidal streams (Binney & Tremaine, 2008, §8.3.1), and it is this case that we investigate in detail in this chapter. Finally, if all the λ_n are zero, there will be no spread at all, and even an unbound cluster will remain intact indefinitely.

We note that there is no *a priori* reason for any of the λ_n to be small. The existence of $\mathbf{D}$ imposes no conditions on H in general, save that it must be twice differentiable near to $\mathbf{J}_0$. We can write a Hamiltonian for which, for particular $\mathbf{J}_0$ at least, the λ_n and the $\hat{\mathbf{e}}_n$ are arbitrary. It must therefore be a peculiar property of realistic Galactic potentials that causes disrupted clusters to form streams.

Lastly, a word on validity. The Taylor expansion equation (5.5) is valid when,

$$D_{ij} \gg \frac{1}{3}\frac{\partial D_{ij}}{\partial J_k}\delta J_k = \frac{1}{3}\frac{\partial^3 H}{\partial J_i \partial J_j \partial J_k}\delta J_k. \tag{5.11}$$

This condition approximates to

$$\delta J_i \ll J_i, \tag{5.12}$$

if H is dominated by some low power of J. In general then, we expect our analysis to be valid if the spread in action of the stars in the cluster is small compared to the actions themselves, which is likely to be true for cold clusters in high-energy orbits around massive hosts. The condition (5.11) is met in detail for all the examples considered below.

5.2.1 The geometry of streams in phase-space

We have seen that the condition for a stream to form is that one of the eigenvalues λ_n of the Hessian, $\mathbf{D}$, must be very much larger than the other two. Herein, we will number the λ_n and their corresponding $\hat{\mathbf{e}}_n$ such that

$$\lambda_1 > \lambda_2 \geq \lambda_3. \tag{5.13}$$

We now ask, under what conditions does the direction of this stream point down the progenitor's orbit?

In action-angle coordinates, the trajectory in angle-space of an orbit $\mathbf{J}$ is given by the frequency vector, $\boldsymbol{\Omega}(\mathbf{J})$. One condition for a stream to delineate precisely the progenitor orbit[1] is that the long axis of the angle-space distribution $\Delta\boldsymbol{\theta}$ must be aligned with the progenitor frequency $\boldsymbol{\Omega}_0$. We will see later that this condition is not sufficient to ensure real-space alignment between streams and orbits, but it is a required condition.

For simplicity, let us restrict ourselves once more to clusters that are isotropic in $\Delta\mathbf{J}$. The angle-space distribution $\Delta\boldsymbol{\theta}$ is now an ellipsoid with semi-axes of length λ_n, and with the semi-major axis of the distribution aligned with the principal direction of $\mathbf{D}$, i.e. $\hat{\mathbf{e}}_1$. The $\hat{\mathbf{e}}_n$ are given by the expression

$$(\nabla_{\mathbf{J}}\boldsymbol{\Omega})\,\hat{\mathbf{e}}_n = \lambda_n\hat{\mathbf{e}}_n, \qquad (n = 1, 3), \tag{5.14}$$

where subscript n denotes a label, i.e. no summation. Since the $\hat{\mathbf{e}}_n$ and λ_n are properties of $\mathbf{D}$ which is assumed constant with respect to $\mathbf{J}$ over the cluster, λ_n and $\hat{\mathbf{e}}_n$ are constant with respect to J. We can therefore rearrange and integrate the above expression

$$\nabla_{\mathbf{J}}(\boldsymbol{\Omega}\cdot\hat{\mathbf{e}}_n) = \lambda_n\nabla_{\mathbf{J}}(\delta\mathbf{J}\cdot\hat{\mathbf{e}}_n), \tag{5.15}$$

$$\boldsymbol{\Omega}\cdot\hat{\mathbf{e}}_n = \lambda_n\,\delta\mathbf{J}\cdot\hat{\mathbf{e}}_n + k_n, \tag{5.16}$$

[1] Alternately, one can consider this orbit to be the mean orbit of the stream stars. Although the actions $\mathbf{J}$ of stars bound to the cluster are not constant, the mean action $\mathbf{J}_0$ does remain constant, since by energy conservation

$$\mathrm{d}H = \boldsymbol{\Omega}_0\cdot\mathrm{d}\mathbf{J}_a + \boldsymbol{\Omega}_0\cdot\mathrm{d}\mathbf{J}_b = 0,$$
$$\mathrm{d}\mathbf{J}_a = -\mathrm{d}\mathbf{J}_b.$$

This restatement of the law of conservation of momentum tells us that any change in the action $\mathbf{J}_a$ of a single star must be matched by an equal and opposite change in the mean action $\mathbf{J}_b$ of the other stars.

where k_n is a constant of integration. Comparing equation (5.16) with equation (5.8) we find

$$k_n = \mathbf{\Omega}_0 \cdot \hat{\mathbf{e}}_n. \tag{5.17}$$

Since $\mathbf{D}$ is symmetric, the $\hat{\mathbf{e}}_n$ are orthogonal. For $\hat{\mathbf{e}}_1$ to be perfectly aligned with $\mathbf{\Omega}_0$, both $\hat{\mathbf{e}}_2$ and $\hat{\mathbf{e}}_3$ must be orthogonal to $\mathbf{\Omega}_0$. Hence

$$k_n = \mathbf{\Omega}_0 \cdot \hat{\mathbf{e}}_n = 0, \qquad (n = 2, 3), \tag{5.18}$$

is the required condition for a stream formed from an isotropic cluster to delineate the progenitor orbit.

Nothing we have said thus far shows that this ought to be the case for either general or specific potentials. In the section that follows, we will determine in detail the form that a Hamiltonian must take, if it is to satisfy this condition.

5.3 Stream formation in spherical potentials

5.3.1 The general case in systems with two actions

In this section, we will explicitly solve equation (5.16) in the case of a general Hamiltonian described by two actions, $H(J_1, J_2)$, such as can be used to describe regular motion in any spherical potential. The resulting solution will be the form of the Hamiltonian in any stream-forming system, and our goal is to relate the terms in that solution to the geometry of stream formation described by $\mathbf{D}$. We will then use these relations to examine the geometry of streams that form in some example systems.

We begin by writing out equation (5.16) explicitly

$$\Omega_1\, e_{n,1} + \Omega_2\, e_{n,2} = e_{n,1}\, \partial_1 H + e_{n,2}\, \partial_2 H = \lambda_n (e_{n,1}\, \delta J_1 + \mathrm{e}_{n,2}\, \delta J_2) + k_n, \tag{5.19}$$

where the $\{1, 2\}$ suffixes denote the respective vector components, and where we have used the shorthand $\partial_i \equiv \partial/\partial J_i$. Equation (5.19) has two solutions, one corresponding to each principal direction $\hat{\mathbf{e}}_n$, where $n = (1, 2)$. We first define the constants

$$\alpha_n \equiv e_{n,2}/e_{n,1}, \tag{5.20}$$

$$\beta_n \equiv k_n/e_{n,1}. \tag{5.21}$$

Equation (5.19) then becomes

$$\partial_1 H + \alpha_n \partial_2 H = \lambda_n(\delta J_1 + \alpha_n \delta J_2) + \beta_n, \tag{5.22}$$

which is a PDE for H in $\delta\mathbf{J}$. The solution to the homogeneous equation for (5.22) is

$$H = f(\delta J_1 - \frac{\delta J_2}{\alpha_n}), \tag{5.23}$$

where f is an arbitrary function of the characteristic coordinate, $\left(\delta J_1 - \frac{\delta J_2}{\alpha_n}\right)$. The general solution for the inhomogeneous equation (5.22) is

$$H = f\left(\delta J_1 - \frac{\delta J_2}{\alpha_n}\right) + \frac{\lambda_n}{2}(\delta J_1^2 + \delta J_2^2) + \beta_n \delta J_1 + H_0. \tag{5.24}$$

This solution relates the form of the Hamiltonian to the geometry of streams that are created within it, with this geometry being explicitly described by $(\alpha, \beta, \lambda)_n$. We must take care to note that although $n = (1, 2)$ and it appears that we have specified two forms for the Hamiltonian, this is not strictly the case. The Hamiltonian has *one* unique form $H(\mathbf{J})$, but the Taylor expansion equation (5.5) and the above equation (5.24) can *always* be compared in two different ways. This can be understood by considering a second-order Taylor expansion of equation (5.24) itself: the quadratic term that would appear in α_n as a result of expanding out f can always be solved for two roots, (α_1, α_2), for the same values of H and $\mathbf{J}$. Indeed, we will perform this calculation explicitly below.

Consider again equation (5.24). There are two possibilities. Firstly, the Hamiltonian of the system may globally take a single form of this type. In this case, the geometry of stream propagation in action-angle space will be the same for all streams of equivalent $\mathbf{\Delta J}$ configuration, no matter what the orbit of the progenitor. We shall see that this is the case with the Kepler potential, discussed below.

The other possibility is that the Hamiltonian does not take a form that can be globally described by equation (5.24) at all. Expanding f in a power series,

$$f(x) = \gamma_n x + \delta_n x^2, \tag{5.25}$$

where we have considered any constant term already subsumed into H_0, and we have neglected any terms

in $\delta\mathbf{J}^3$, gives

$$\begin{aligned} H &= H_0 + \gamma_n \left(\delta J_1 - \frac{\delta J_2}{\alpha_n}\right) + \delta_n \left(\delta J_1 - \frac{\delta J_2}{\alpha_n}\right)^2 + \frac{\lambda_n}{2}\left(\delta J_1^2 + \delta J_2^2\right) + \beta_n \delta J_1 \\ &= \delta J_1^2 \left(\delta_n + \frac{\lambda_n}{2}\right) + \delta J_2^2 \left(\frac{\delta_n}{\alpha_n^2} + \frac{\lambda_n}{2}\right) - 2\frac{\delta_n}{\alpha_n}\,\delta J_1\,\delta J_2 + \delta J_1\left(\gamma_n + \beta_n\right) - \frac{\gamma_n}{\alpha_n}\delta J_2 + H_0. \end{aligned} \tag{5.26}$$

We rewrite the Taylor expansion (5.5) in our current notation,

$$H(\mathbf{J}) = H_0 + \Omega_{0,i}\,\delta J_i + \frac{1}{2} D_{ij}\,\delta J_i\,\delta J_j. \tag{5.27}$$

Comparing coefficients between the above expression and equation (5.26) gives,

$$\begin{aligned} \alpha_n &= \frac{-2\delta_n}{\partial_1 \Omega_{0,2}}, \\ \beta_n &= \Omega_{0,1} + \alpha_n \Omega_{0,2}, \\ \gamma_n &= -\alpha_n \Omega_{0,2}, \\ \delta_n &= \frac{1}{4}\left\{(\partial_1\Omega_{0,1} - \partial_2\Omega_{0,2}) \pm \sqrt{(\partial_1\Omega_{0,1} - \partial_2\Omega_{0,2})^2 + 4(\partial_1\Omega_{0,2})^2}\right\}, \\ \lambda_n &= \partial_1\Omega_{0,1} - 2\delta_n. \end{aligned} \tag{5.28}$$

where we can now see explicitly that choosing $n = (1, 2)$ is equivalent to setting the sign of the radical in the above expression for δ_n.

We are now in a position to deduce the geometry of streams for a given potential, if we know the form of $H(\mathbf{J})$. In particular, we can deduce the angle between the principal direction $\hat{\mathbf{e}}_1$ and $\mathbf{\Omega}_0$, since

$$\mathbf{\Omega}_0 = \left\{\gamma_n + \beta_n, \frac{-\gamma_n}{\alpha_n}\right\}, \tag{5.29}$$

and from equation (5.17) and the definition of β_n (5.21),

$$\mathbf{\Omega}_0 \cdot \hat{\mathbf{e}}_1 = e_{1,1}\beta_1, \tag{5.30}$$

and so the misalignment angle is

$$\vartheta = \arccos \frac{e_{1,1}\beta_1}{|\mathbf{\Omega}_0|}. \tag{5.31}$$

Table 5.1: Parameters for the spherical potentials used in this chapter.

	$M/10^{10}M_\odot$	$b/\,\mathrm{kpc}$
Kepler	10.75	0
Isochrone	28.52	3.64

where, since the $\hat{\mathbf{e}}_n$ are normalized to unity

$$e_{1,1} = \pm\sqrt{\frac{1}{1+\alpha_1^2}}. \tag{5.32}$$

We will now proceed to study the formation of streams in various example potentials.

5.3.2 Kepler potential

In the study of galactic dynamics, there are remarkably few potentials of interest for which a Hamiltonian can be written as a function of $\mathbf{J}$ in closed form. One such potential is the Kepler potential

$$\Phi(r) = \frac{-GM}{r}, \tag{5.33}$$

for which the Hamiltonian is,

$$H(\mathbf{J}) = \frac{-(GM)^2}{2(J_r + L)}, \tag{5.34}$$

where J_r is the radial action, and L is the angular momentum. Comparing the above expression with equation (5.24), we immediately see that H is of the form required for streams to form with globally consistent geometry. The most obvious solution for the parameters (5.28) is

$$(\alpha, \beta, \lambda)_2 = (-1,\, 0,\, 0)\,, \tag{5.35}$$

for $\hat{\mathbf{e}}_2$, which has a null eigenvalue, and lies perpendicular to $\mathbf{\Omega}_0$. Some algebra will also confirm the other solution

$$(\alpha, \beta, \lambda)_1 = (1,\, 2\Omega_0,\, 2\partial_{J_r}\Omega_0)\,, \tag{5.36}$$

for the principal direction $\hat{\mathbf{e}}_1$, which points precisely along $\mathbf{\Omega}_0$. Streams therefore perfectly delineate progenitor orbits in Kepler potentials, and exhibit secular spread strictly along the orbit, and do not grow wider with time.

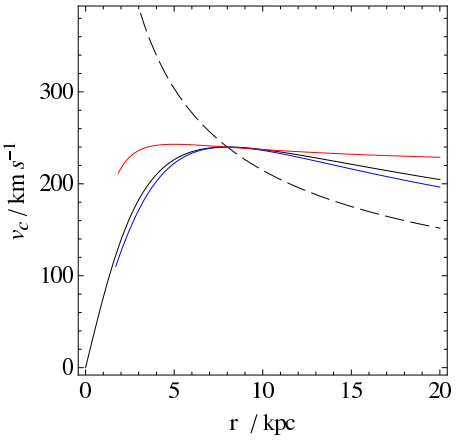

Figure 5.1: The in-plane rotation curves of the potentials specified in Table 5.1 and Table 5.5. Solid black line: the isochrone potential. Dashed black line: the Kepler potential. Red line: the Stäckel potential SP2. Blue line: the Stäckel potential SP1. Parameters in all potentials have been tuned to give a circular speed of $v_c = 240\,\mathrm{km\,s^{-1}}$ at the Solar radius $R_0 = 8\,\mathrm{kpc}$.

Table 5.2: Actions and apses for selected orbits in the spherical potentials used in this chapter. We generally choose $J_\phi = L$ so all orbits remain in the (x, y) plane. The trajectories of these orbits are illustrated in Fig. 5.2

	J_r/ kpc km s^{-1}	L/ kpc km s^{-1}	r_p/ kpc	r_a/ kpc
K1	780.	1016.	1.5	13
I1	313.	1693.	5	13
I2	0.	1920.	8	8
I3	207.	1920.	6	12.5
I4	571.7	2536.	11	20
I5	215.4	3127.	7	20

5.3.2.1 Numerical tests

Table 5.1 describes a Kepler potential in which we will now confirm this prediction numerically. The parameter M of this potential was chosen to reproduce a fiducial circular velocity $v_c = 240\,\mathrm{km\,s^{-1}}$ at the approximate solar radius of $R_0 = 8\,\mathrm{kpc}$. Fig. 5.1 shows the rotation curve in this potential, along with the rotation curves for other potentials in use in this chapter.

A cluster of 50 test particles was created, with a Gaussian distribution of particles in action-angle space, defined by $\sigma_J = 1\,\mathrm{kpc\,km\,s^{-1}}$ and $\sigma_\theta = 5 \times 10^{-3}$ radians. This cluster was placed on the orbit K1, given in Table 5.2. in the aforementioned Kepler potential. The orbit has a pericentre radius of about 1.5 kpc and an apocentre of about 13 kpc, and is illustrated in Fig. 5.2.

The cluster was released at apocentre, and evolved for 4.02 Gyr, equal to 24 complete orbits, by integrating the equations of motion for each particle in the relevant potential. Fig. 5.3 shows the real-space configuration of the particles at the end of this time, and Fig. 5.4 shows the configuration of the same particles in angle-space.

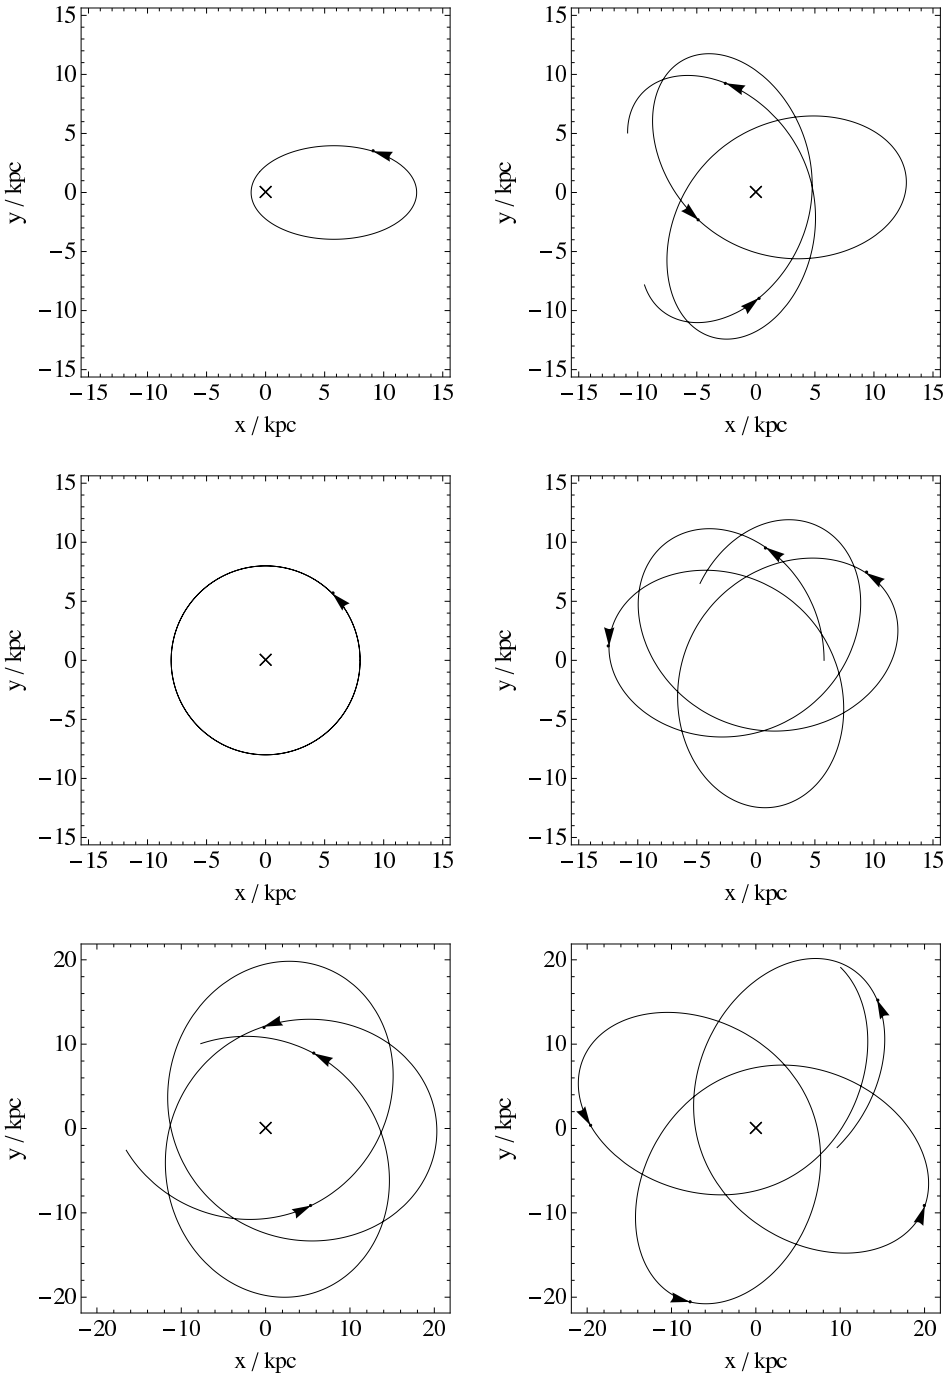

Figure 5.2: The real-space trajectories of the orbits (left-to-right, top-to-bottom) K1, I1, I2, I3, I4 and I5, as described in Table 5.2. The trajectories were evaluated in the Kepler potential (for K1) and the isochrone potential (for I1–I5) of Table 5.1.

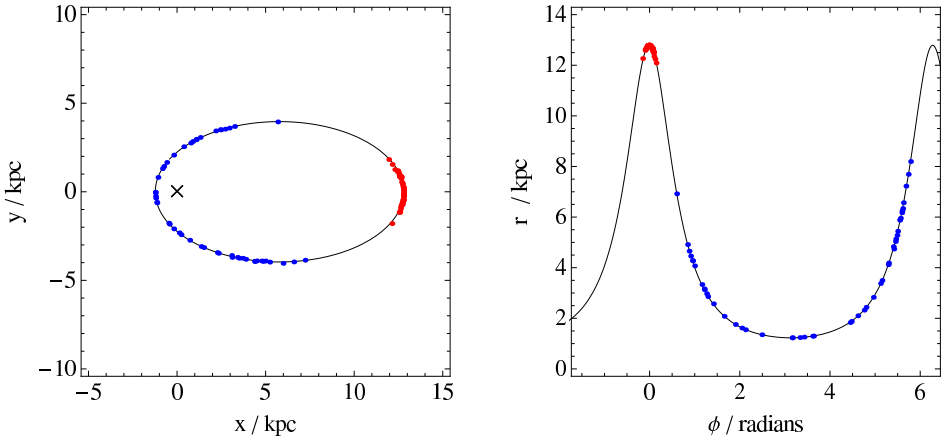

Figure 5.3: The solid line shows the orbit K1 (Table 5.2) in a Kepler potential (Table 5.1), on which a cluster of 50 test particles (initial conditions described in the text) has been evolved. The particles were released at apocentre. The red dots show the position of the test particles near apocentre, after 24 complete orbits, at $t = 4.02\,\mathrm{Gyr}$. The blue dots show the same test particles near pericentre, approximately half an orbit later. In both cases, the dots delineate the underlying orbit precisely.

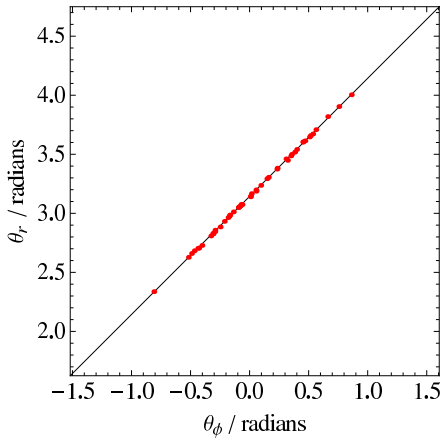

Figure 5.4: Angle-space configuration for the particles shown at apocentre in Fig. 5.3. The solid line shows the frequency vector $\mathbf{\Omega}_0$, with which the stream particles are perfectly aligned. There is also no secular spread in the direction perpendicular to the stream motion, as predicted in §5.3.2.

The cluster has elongated to form a stream, that Fig. 5.3 shows to cover approximately half the orbit when the centroid is at pericentre. There is no spread in width in either real-space or angle-space. The stream delineates the cluster's orbit in both figures perfectly. Fig. 5.3 shows that this is true irrespective of the real-space location of the centroid. Thus, the prediction of the previous section is validated.

5.3.3 Spherical harmonic oscillator

The spherical harmonic oscillator potential applies for motion within a sphere of uniform density (Binney & Tremaine, 2008, §3.1a), and is therefore of relevance to galaxy cores in the absence of a black hole. It has the form

$$\Phi(r) = \frac{1}{2}\Omega^2 r^2. \tag{5.37}$$

The Hamiltonian of the harmonic oscillator takes a particularly simple form

$$H(J_r, L) = \Omega\left(L + 2J_r\right), \tag{5.38}$$

where Ω is a constant. Comparing the above expression with equation (5.24), we see that like the Kepler Hamiltonian, the form of equation (5.38) admits the same stream geometry everywhere. In this case, however, the solutions to equation (5.28) are trivial, since $\mathbf{D}$ is a null matrix, and β_n and λ_n are identically zero for both $n = (1, 2)$.

We conclude that clusters in harmonic potentials will always remain in the same configuration in angle-space, and will not spread out. Hence, streams cannot form in harmonic potentials. We further note that, in any case, it would be difficult to tidally strip a cluster in a harmonic potential, since the tidal force $dF_{\rm tide}$ across a cluster

$$dF_{\rm tide} \simeq \frac{\partial^2\Phi}{\partial r^2}\mathrm{d}r = \Omega^2\mathrm{d}r, \tag{5.39}$$

is independent of galactocentric radius r. Thus, a cluster that is bound at apocentre in such a potential will remain bound elsewhere along its orbit.

5.3.4 Isochrone potential

The isochrone potential (Binney & Tremaine, 2008, §2.2.2d) is a simple potential which has several useful properties. It behaves as a harmonic oscillator in the limit of small radius, and as a Kepler potential at large radius, thus providing a reasonable model for a spherical galaxy across all radii. The form of the

potential is

$$\Phi(r) = \frac{-GM}{b + \sqrt{b^2 + r^2}}, \tag{5.40}$$

where b is a scaling constant, and the Hamiltonian is

$$H(\mathbf{J}) = \frac{-(GM)^2}{2[J_r + \frac{1}{2}(L + \sqrt{L^2 + 4GMb})]^2}. \tag{5.41}$$

This Hamiltonian is not of the form of equation (5.24), and does not admit a globally applicable stream geometry. We therefore need to calculate the parameters (5.28) directly from H. We require the frequencies, which are obtained by direct differentiation

$$\Omega_r = \frac{(GM)^2}{[J_r + \frac{1}{2}(L + \sqrt{L^2 + 4GMb})]^3}, \tag{5.42}$$

$$\Omega_\phi = \frac{1}{2}\left(1 + \frac{L}{\sqrt{L^2 + 4GMb}}\right)\Omega_r. \tag{5.43}$$

We also require the derivatives of the frequencies with respect to the actions. We note that Ω_r can be written as a function of H only,

$$\Omega_r = \frac{(-2H)^{3/2}}{GM}, \tag{5.44}$$

and therefore its derivatives with respect to the actions are

$$\partial_{J_r}\Omega_r(H) = \Omega_r'(H)\,\Omega_r, \tag{5.45}$$

$$\partial_L\Omega_r(H) = \Omega_r'(H)\,\Omega_\theta, \tag{5.46}$$

where

$$\Omega_r'(H) = -\frac{3\sqrt{-2H}}{GM}. \tag{5.47}$$

We must calculate $\partial_L\Omega_\phi$ directly from equation (5.43). We find

$$\partial_L\Omega_\phi = \frac{2GMb}{(L^2 + 4GMb)^{3/2}}\Omega_r + \frac{1}{2}\left(1 + \frac{L}{\sqrt{L^2 + 4GMb}}\right)\partial_L\Omega_r. \tag{5.48}$$

From these expressions, we are now in a position to piece together values for the parameters (5.26), although the full expressions for $(\alpha, \beta, \gamma, \delta, \lambda)_n$ are not algebraically neat, and are therefore not very instructive. We do not repeat them here. We do note however that in the limit of $b \to \infty$ we recover the

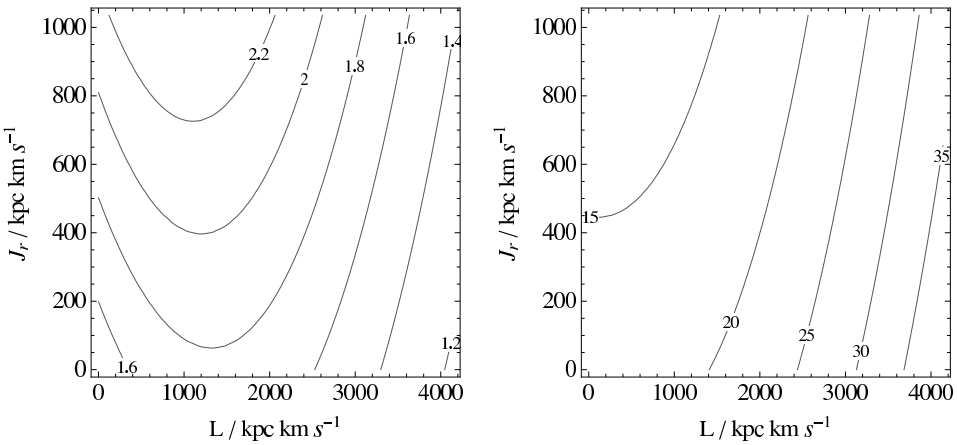

Figure 5.5: Details of stream geometry in the isochrone potential of Table 5.1. Left panel: the misalignment angle ϑ, in degrees, between the principal direction of $\mathbf{D}$ and $\mathbf{\Omega}_0$, shown as a contour plot against the actions of the progenitor orbit. In all cases, ϑ is between 1.2° and 2.2° in angle-space. Right panel: The ratio of the eigenvalues λ_1/λ_2. The ratio is > 10 everywhere and rises sharply with increasing L. The actions shown in both plots cover a range of interesting orbits, which are described in Table 5.3.

Table 5.3: The coordinate extrema of selected orbits from Fig. 5.5, illustrating the variety of orbits covered by that figure. The actions are expressed in $\mathrm{kpc\,km\,s^{-1}}$, while the apses are in kpc.

J_r	L	r_p	r_a
1000	4000	13	46
1000	~ 0	~ 0	12
0	4000	20.3	20.3

harmonic oscillator case of §5.3.3, and that in the limit of $b \to 0$ we recover the Kepler case of §5.3.2, as is required.

In order to proceed, we must work with a specific example. Table 5.1 describes an isochrone potential, chosen to have the rotation curve maximized at $v_c = 240\,\mathrm{km\,s^{-1}}$ at the assumed solar radius of $R_0 = 8\,\mathrm{kpc}$. The rotation curve for this potential is plotted in Fig. 5.1.

What then is the geometry of streams formed in this potential? Fig. 5.5 shows the misalignment angle ϑ, given by equation (5.31), and the ratio of the eigenvalues λ_1/λ_2, both as functions of $\mathbf{J}$. The range of $\mathbf{J}$ shown covers a variety of interesting orbits, described in Table 5.3.

The left panel shows that the principal direction of $\mathbf{D}$ is misaligned with the progenitor orbit for all values of $\mathbf{J}$, by 1–2°. The misalignment is at a minimum for both low energy and high energy circular orbits, and at a maximum for eccentric orbits with a guiding centre close to $r = b$. The right panel shows that the ratio of the eigenvalues λ_1/λ_2 varies from 15 to 25 across the range, with the ratio maximized for high energy circular orbits, and minimized for high energy plunging orbits.

Thus, we expect an isotropic cluster of test particles in this potential to form a stream in angle-space that is misaligned with $\mathbf{\Omega}_0$ by 1–2° and is 15–20 times longer than it is wide.

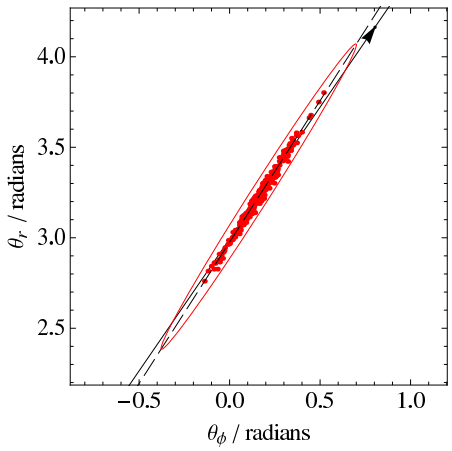

Figure 5.6: The angle-space distribution of a cluster of test particles, evolved on orbit I1 (Table 5.2) in the isochrone potential of Table 5.1 for 94 complete azimuthal circulations. The particles are shown at apocentre. The frequency vector $\mathbf{\Omega}_0$ of the progenitor orbit is shown with an arrowed black line. The stream is slightly misaligned with $\mathbf{\Omega}_0$; the dashed line is a straight line fit to the positions of the test particles, and clearly demonstrates this misalignment. Also plotted with a red solid line is the image in angle-space of a circle in action-space, mapped by $\mathbf{D}$. The shape and orientation of the image reflects the λ_n and $\hat{\mathbf{e}}_n$ of $\mathbf{D}$ for this orbit. The ellipse is clearly misaligned with the underlying orbit, but is perfectly aligned with the stream particles.

5.3.4.1 Numerical tests

We have created a cluster of 150 test particles, randomly sampled from a Gaussian distribution in action-angle space, defined by $\sigma_J = 0.2\,\mathrm{kpc\,km\,s^{-1}}$ and $\sigma_\theta = 10^{-3}$ radians. This cluster was placed on the orbit I1 from Table 5.2, which has an apocentre radius of 13 kpc and a pericentre radius of ~ 5 kpc, and is illustrated in Fig. 5.2.

The cluster was released at apocentre, and evolved for 94 complete azimuthal circulations, equal to a period of $t = 22.75$ Gyr. Fig. 5.6 shows the angle-space configuration of the particles after this time. The arrowed, black line shows the orbit of the underlying cluster, $\mathbf{\Omega}_0$. The dashed line shows a straight line fit to the distribution of particles, which is clearly misaligned with the black line. We further note that, unlike the stream in the Kepler potential shown in Fig. 5.4, this stream is clearly also increasing in width.

We can predict the shape of this distribution precisely. Plotted as a red ellipse in Fig. 5.6 is the angle-space image of a circle in action-space, having been mapped by $\mathbf{D}$. After some time, the angle-space distribution of an isotropic cluster of test particles should take the form of a scaled version of this image. We see that the image and the particle distribution are indeed comparable, and that the dashed line is perfectly aligned with the principal axis of the image.

How does this misalignment manifest itself in real-space? Fig. 5.7 shows the real-space configuration of the cluster at the end of the simulation. The left panel shows an overview of the cluster in the orbital

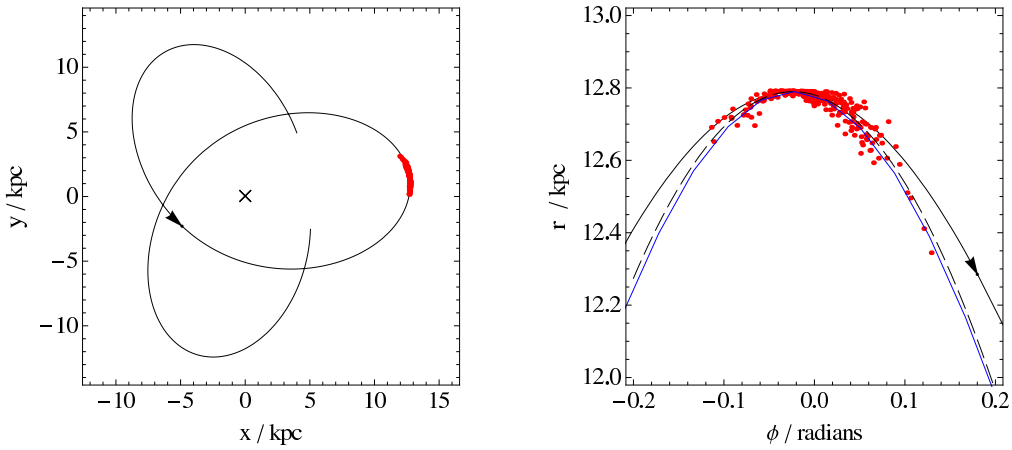

Figure 5.7: Real-space configuration for the stream of test particles shown in Fig. 5.6. The arrowed black line shows the trajectory of the progenitor orbit. The left panel shows an overview of the particles and the orbit; the cross marks the centre of the potential. The right panel shows a zoomed-in view of the particles, and is plotted in polar coordinates. The stream formed by the dots falls away in radius faster than does the orbit, in both forwards and backwards directions: the stream clearly does not follow the progenitor orbit. The dashed line is a quadratic curve least-squares fitted to the stream, which shows that the stream has a substantially greater curvature than does the underlying orbit. The solid blue line is the track predicted in §5.4.3 from the dashed line in Fig. 5.6: it agrees perfectly with the stream.

plane, while the right panel shows a close-up view of the cluster. The orbit of the cluster is drawn with a solid black line. In the right panel, the particles have been least-squares fitted to a quadratic curve, shown as a dashed line.

The progenitor orbit is clearly a poor representation of the stream. Although the orbit passes through the centroid of the stream, as expected, the curvature of the orbit is too low to match the stream adequately. Thus, the small misalignment in angle-space is manifest as a significant change in stream curvature at apocentre.

We conclude that, in this case, the track of the stream makes a poor proxy for the orbit of its stars, and that in general, streams cannot be relied upon to delineate orbits. In §5.5 we will highlight the errors that can be made in attempting to optimize potential parameters based on the assumption that they do delineate orbits. Firstly, however, we will discuss the details involved in the mapping of streams from action-angle space to real-space, in order that we may properly predict the track of a stream.

5.4 Mapping streams from action-angle space to real space

5.4.1 Non-isotropic clusters

Up until this point, we have considered our streams to form from a cluster of particles that is isotropic in $\mathbf{J}$, resulting in a stream that is perfectly aligned in angle-space with the principal direction of $\mathbf{D}$.

It is not obvious that this is a fair assumption. The structure in angle-space, given by equations (5.9) and (5.10), is linearly dependent upon the action-space distribution that generates it. By properly choosing that distribution, we can create streams of arbitrary shape in angle-space[2]. Clearly, nature does not create clusters with arbitrary action-space distributions, so arbitrary-shaped streams do not emerge. But what kind of action-space distribution should be considered reasonable, that we may think of how it maps?

In §5.6 below, we will investigate the action-space distribution of real clusters, by means of N-body simulation. Here, we only require a qualitative understanding, in order to guess what kind of action-angle structures we should learn to map. Since the action-space distribution arises from the random motion of stars within the cluster, there is unlikely to be much complex structure. We will therefore assume that the distribution is ellipsoidal. But what should be the axis ratio of this ellipse?

Consider a cluster with isotropic velocity dispersion, σ, that is on an orbit with apocentre r_a and pericentre r_p, where it is tidally disrupted. In our spherical system, the radial action is given by the closed integral

$$J_r = \frac{1}{2\pi} \oint p_r \, dr, \tag{5.49}$$

where the integration path is one complete radial oscillation. Now consider a particle whose radial momentum p_r differs from that of the cluster average by $\delta p_r \sim \sigma$. We can take a finite difference over equation (5.49) and thus obtain an expression for the the difference in radial action between the particle and the cluster

$$\delta J_r \sim \frac{1}{\pi} \delta p_r \Delta r \sim \sigma \Delta r, \tag{5.50}$$

where $\Delta r = (r_a - r_p)$ is the amplitude of the radial oscillation. Now consider another particle, whose azimuthal velocity differs from that of the cluster by $\delta v_t \sim \sigma$. The difference in angular momentum between this particle and the cluster is

$$\delta L \sim r_p \, \delta v_t \sim r_p \sigma, \tag{5.51}$$

where we have performed our calculation at pericentre, because that is where the cluster is stripped. For the purposes of this section, we are interested in the relative size of the spread in radial action ΔJ_r and

[2]Systems in which one of the λ_i is null are the exception to this statement. Structures in such systems are limited to that subset of angle-space which is spanned by the non-null eigenvectors of $\mathbf{D}$. The Kepler potential is one such system; all structures are mapped to a line pointing precisely along the frequency vector $\mathbf{\Omega}_0$.

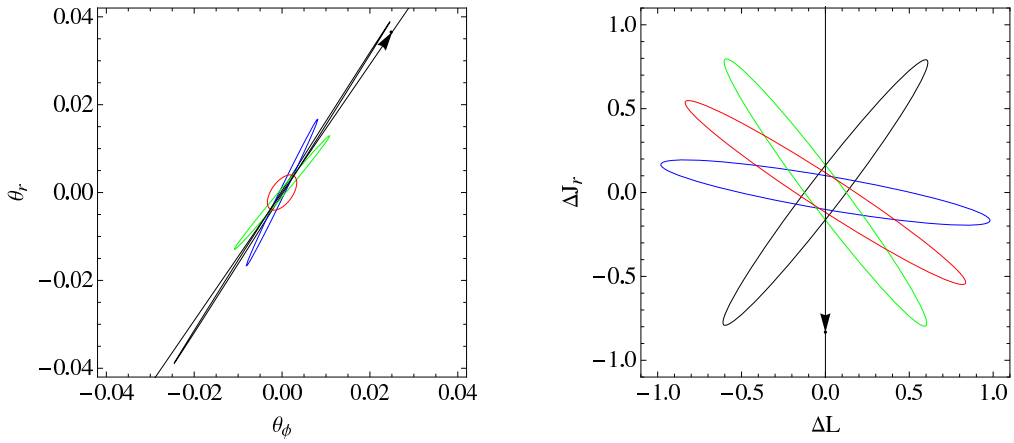

Figure 5.8: The right panel shows a selection of ellipses in action-space, each of axis ratio 10, but oriented in different directions. The left panel shows the image in angle-space that results from mapping each of the action-space ellipses with the Hessian $\mathbf{D}$, calculated for the orbit I1 in the isochrone potential of Table 5.1. In the left panel, the arrowed black line is the frequency vector $\mathbf{\Omega}_0$; in the right panel, the arrowed black line is the inverse map of the frequency vector, $\mathbf{D}^{-1}\mathbf{\Omega}_0$. We see that regardless of the shape in action-space, the mapped images are all elongated and roughly aligned with the orbit, although the alignment is generally not perfect. In this example, the misalignment of the red and green images is about 10°.

the spread in angular momentum ΔL for a disrupting cluster. We see that is approximately given by

$$\frac{\Delta J_r}{\Delta L} \sim \frac{\Delta r}{\pi r_p}. \tag{5.52}$$

Although this ratio will take on every value between $(0, \infty)$ as we move from a circular orbit to a plunging one, for the orbits likely to be occupied by stream-forming clusters, it will typically be of order unity.

The right panel of Fig. 5.8 shows a set of ellipses, with axis ratio 10, placed at various orientations in action-space. The left panel of Fig. 5.8 shows the images of these ellipses in angle-space, following mapping by $\mathbf{D}$ when evaluated on I1. Note that all the images are both elongated and roughly oriented towards the principal direction. We conclude that the images of most action-space ellipses under this map—and thus, most streams formed in this potential—would be highly elongated and oriented to within a few degrees of the principal direction, which is itself oriented to within a few degrees of the frequency vector $\mathbf{\Omega}_0$.

Since the ratio of the eigenvalues for this orbit is ~ 17, it is not possible to produce an image in angle-space that is not elongated towards the principal direction, by mapping an action-space ellipse of axis-ratio 10. We note from Fig. 5.5 that, for this potential, the ratio of eigenvalues does not vary much, and nor does the principal direction stray from $\mathbf{\Omega}_0$ by more than a few degrees. We therefore conclude that reasonable action-space distributions will always result in the formation of streams in this potential, and that such streams will always be oriented in angle-space to within a few degrees of $\mathbf{\Omega}_0$.

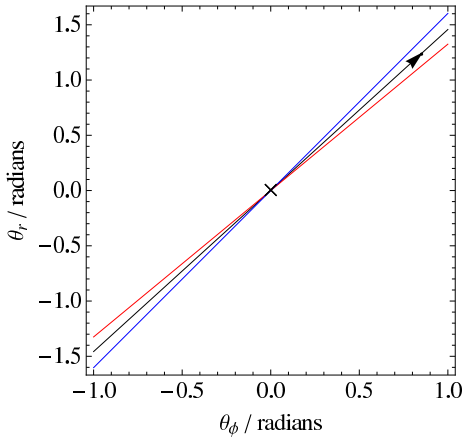

Figure 5.9: Three trajectories in angle-space. The black line is the trajectory of orbit I1, in the isochrone potential of Table 5.1. The red line has a frequency ratio Ω_r/Ω_ϕ that is 10 per cent lower than I1. Conversely, the blue line has a frequency ratio that is 10 per cent higher than I1. We note that the red line has retarded radial phase (relative to the black line) on the leading tail, and advanced radial phase on the trailing tail. Conversely, the blue line has advanced radial phase in the leading part, and retarded radial phase in the trailing part.

5.4.2 The mapping of action-angle space to real space

Fig. 5.9 shows three trajectories in angle space. The black line is the trajectory of I1 in the isochrone potential given in Table 5.1. The red line has a frequency ratio Ω_r/Ω_ϕ that is 10 per cent lower than the black line, and is therefore rotated from it by approximately 2.9°. Conversely, the blue line has a frequency ratio that is 10 per cent higher than the black line, and is therefore rotated from it by about 2.5°.

The red and blue lines were chosen to represent likely streams in angle-space that could form in the isochrone potential, given the results of the previous section. We note that the red line has retarded radial phase (relative to the black line) on the leading tail, and advanced radial phase on the trailing tail. Conversely, the blue line has advanced radial phase in the leading part, and retarded radial phase in the trailing part.

How do these lines map into real-space? Fig. 5.10 shows the real-space curve obtained from the lines in Fig. 5.9, having chosen the point of intersection such that the lines are phase-matched near apocentre. The mapping is done by solving numerically for the real-space roots of the equations that relate action-angle variables to real-space coordinates: for spherical potentials, the appropriate equations are given in §3.5.2 of Binney & Tremaine (2008). The curves in Fig. 5.10 have been drawn by assuming that all points along each line have the same $\mathbf{J}$. Thus, this figure represents the real-space curves of streams oriented in angle-space according to Fig. 5.9, but formed from clusters of vanishingly small $\Delta\mathbf{J}$.

In Fig. 5.10 we see that the red line has systematically lower curvature than the black line. Conversely,

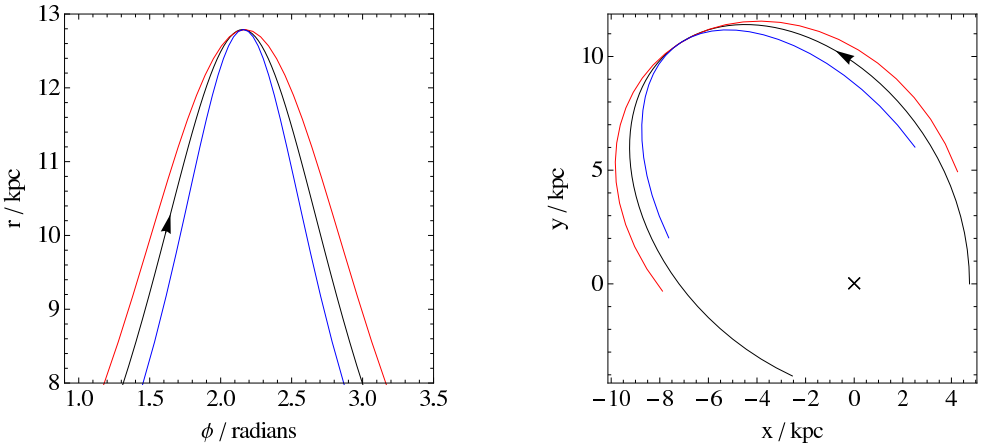

Figure 5.10: Plots of the real-space trajectories of the lines shown in Fig. 5.9, phase-matched near apocentre. Left panel: the trajectories plotted in polar coordinates. Right panel: the trajectories plotted in Cartesian coordinates; the centroid of the potential is marked with a cross.

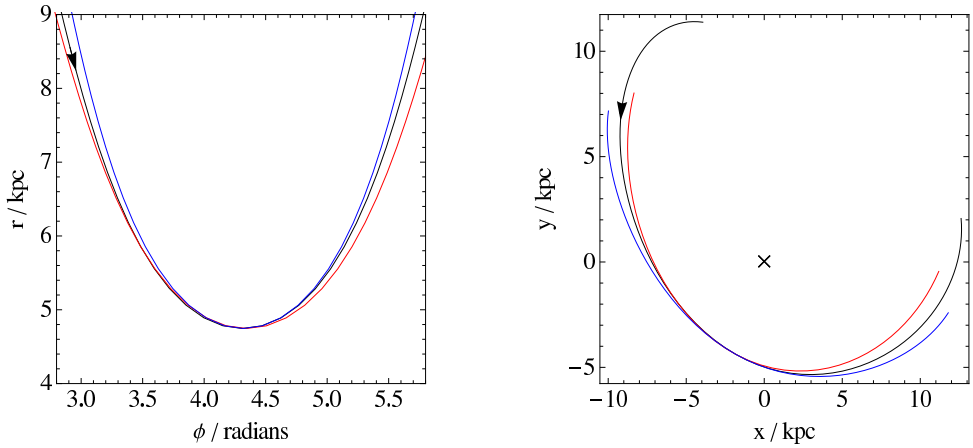

Figure 5.11: Similar to Fig. 5.10, but with the trajectories phase-matched near pericentre. The red line is again observed to have systematically lower curvature than the black line, and the blue line is again observed to have systematically greater curvature than the black line.

the blue line has systematically greater curvature than the black line. We understand this, because the red curve has retarded radial phase on the leading tail, and advanced radial phase on the trailing tail, and thus is flattened with respect to the orbit. Similarly, the blue line has advanced radial phase on the leading tail, and retarded radial phase on the trailing tail, and thus appears curved with respect to the orbit.

Fig. 5.11 shows the same lines, but now phase-matched at pericentre. Similarly to Fig. 5.10, the red line again appears flattened with respect to the orbit, and the blue line appears curved with respect to the orbit. Fig. 5.12 also shows the same lines phase-matched at a point well away from apsis. In this case, a misalignment between the stream and $\mathbf{\Omega}_0$ in angle-space is expressed as a real-space misalignment, rather than a curvature error as at apsis. We note the similarity between the left panel and Fig. 5.9

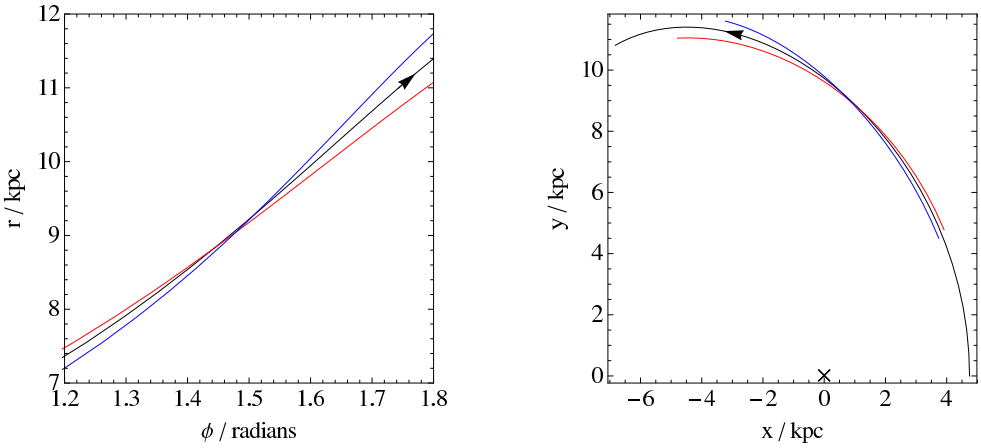

Figure 5.12: Similar to Fig. 5.10, but with the trajectories phase-matched at a point well away from apsis. In this case, a misalignment between the stream and $\mathbf{\Omega}_0$ in angle-space is expressed as a real-space misalignment, rather than a curvature error as at apsis. We note the similarity between the left panel and Fig. 5.9.

which occurs because, unlike at apsis, the mapping between angle-space and plane polar coordinates is relatively undistorted near this point.

5.4.3 Numerical tests

How can we be sure that our calculation of stream tracks from lines in angle-space is correct? We constructed a cluster of 100 test particles by randomly sampling a segment of a line in action-space, chosen such that its angle-space image would be a line oriented perpendicularly to $\mathbf{\Omega}_0$.

A point on the orbit I1 was chosen by integrating backwards 23.6 Gyr from apocentre. The cluster was placed at this point, released, and evolved forwards for 23.6 Gyr. Fig. 5.13 shows the real-space configuration of the resulting stream at the end of this time. Also plotted are the orbit I1, as a black curve, and the stream track predicted from the line-segment in action space, as a red curve. The stream and the prediction are seen to agree almost perfectly[3].

We may further test our apparatus by predicting the track of the stream shown in Fig. 5.7. The blue line in the right panel of that figure shows the stream track predicted from a line in angle-space that is oriented precisely along the principal direction of $\mathbf{D}$ when evaluated for the orbit I1. The blue line is a much better match to both the stream data and the fitted line than is the orbit.

[3]The slight phase mismatch between the apocentre location of the particles and the apocentre location of the predicted track is due to finite numerical precision in the translation between action-angle variables and position-velocity coordinates when calculating the initial conditions for the simulation.

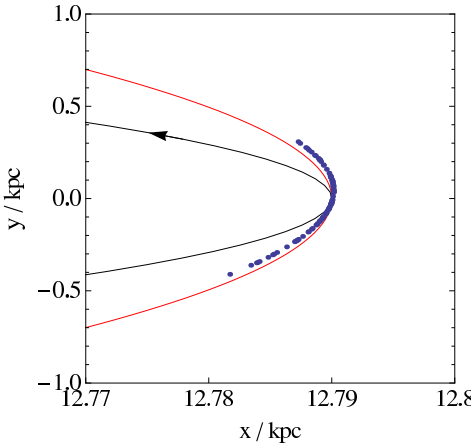

Figure 5.13: Real-space configuration of a 100 particle test stream, intended to check that our mapping of angle-space to real-space is accurate. The stream has an action-space distribution chosen to map under $\mathbf{D}$ to a line in angle-space oriented perpendicularly to $\boldsymbol{\Omega}_0$. The stream was released on orbit I1 in the isochrone potential of Table 5.1 and integrated for 23.6 Gyr. The black curve is the orbit I1, while the red curve is the predicted stream track. The red curve is clearly a far better match to the particles than the black curve; the slight phase mismatch between the apocentre of the particles and the red curve is due to numerical error in the simulation process.

5.4.4 Are trajectories insensitive to small changes in J?

So far, all the real-space tracks we have computed from streams in angle-space have assumed that the corresponding action for all points along that stream is that of the progenitor, $\mathbf{J}_0$.

This assumption is only strictly valid in the case of a vanishingly small action-space distribution, and for asymptotically large time since disruption of the cluster. If a mapping into real-space from a line in angle-space is made under this assumption, then a stream generated by a sufficiently broad action-space distribution will not be accurately represented, even though the representation in angle-space may be exact. This is because the small changes in action that give rise to the small changes in frequency also cause small changes in real-space trajectory as well.

When computing a stream track in real-space, it is possible to correct for this effect. By inverting equation (5.10) and eliminating $\Delta\boldsymbol{\Omega}$ using equation (5.9), we find that for a star separated from a fiducial point on the stream by angle $\delta\boldsymbol{\theta}$, the difference in action between the star and the fiducial point, $\delta\mathbf{J}$ is given by

$$\delta\mathbf{J} = \frac{1}{t_d}\mathbf{D}^{-1}\,\delta\boldsymbol{\theta}, \tag{5.53}$$

where t_d is the time since the star and the fiducial point were coincident. We may therefore guess the correction $\delta\mathbf{J}$ for a star's true action $(\mathbf{J}_0 + \delta\mathbf{J})$ from its position in the stream, provided we know t_d.

We typically take the fiducial point to be the centroid of the stream, following which we may assume

t_d to be the time since the first pericentre passage of the cluster on its present orbit. Although this assumption neglects the possibility that the star could have been torn away during a subsequent pericentre passage, we note that during tidal disruption, it is the fastest moving stars which become unbound. The cluster core that remains after a pericentre passage therefore has lower velocity dispersion. Stars subsequently torn from that cluster will therefore have a smaller distribution in action-space. Consequently, the stars with the largest $\delta\boldsymbol{\theta}$ from the centroid—i.e. those for which the $\delta\mathbf{J}$ correction will be most important—must have been torn away at the earliest time, and so the assumption that t_d equals the time since the first pericentre passage remains good.

But just how important is this effect? For small changes in J_r, the trajectory changes we discuss are expressed as changes in the radial amplitude Δr, while the guiding centre radius r_g, which is purely a function of L, is held constant. We can estimate the magnitude of the effect as follows.

Consider a cluster on an orbit close to circular, whose radial action is given by equation (5.49). Orbital energy E is conserved, so close to apsis $r = r_0$, the radial momentum p_r is given according to

$$E = p_r(r)^2 + \Phi_{\rm eff}(r) = p_r(r)^2 + \Phi_{\rm eff}(r_0) + \left.\frac{\mathrm{d}\Phi_{\rm eff}}{\mathrm{d}r}\right|_{r_0}(r - r_0), \tag{5.54}$$

where we have defined the effective potential $\Phi_{\rm eff}(r) = \Phi(r) + L^2/2r^2$. Since at apsis $p_r = 0$, then

$$E = \Phi_{\rm eff}(r_0), \tag{5.55}$$

so from equation (5.54) we see

$$p_r(r) = \sqrt{-(r - r_0)\left.\frac{\mathrm{d}\Phi_{\rm eff}}{\mathrm{d}r}\right|_{r_0}} = \sqrt{(r - r_0)F_{\rm eff}(r_0)}, \tag{5.56}$$

where we have defined the effective force, $F_{\rm eff}(r') = -\partial\Phi_{\rm eff}/\partial r|_{r'}$. If (r_a, r_p) are apocentre and pericentre respectively, then we see that

$$p_r(r) \simeq \begin{cases} \sqrt{|F_{\rm eff}(r_a)|(r_a - r)} & \text{if } r \simeq r_a, \\ \sqrt{|F_{\rm eff}(r_p)|(r - r_p)} & \text{if } r \simeq r_p. \end{cases} \tag{5.57}$$

Hence, we might define the global approximation to p_r

$$\tilde{p_r}(r) = \frac{\sqrt{|F_{\rm eff}|\,(r - r_p)(r_a - r)}}{\sqrt{r_a - r_p}}, \tag{5.58}$$

with $F_{\rm eff}$ a constant set equal to the value of $F_{\rm eff}(r)$ taken at apsis, at both of which we assume it to

take approximately the same value. We note that

$$\int_{r_p}^{r_a} \sqrt{(r-r_p)(r_a-r)}\,\mathrm{d}r = \frac{\pi}{8}\Delta r^2, \tag{5.59}$$

so that when combined with equation (5.58), equation (5.49) becomes,

$$J_r = \frac{1}{8}\sqrt{|F_{\rm eff}|\,\Delta r^3}. \tag{5.60}$$

We can deduce the value of $F_{\rm eff}$ as follows. We note that

$$\Phi_{\rm eff} = \Phi + \frac{L^2}{2r^2}, \tag{5.61}$$

and that its derivative is

$$\frac{\mathrm{d}\Phi_{\rm eff}}{\mathrm{d}r} = \frac{\mathrm{d}\Phi}{\mathrm{d}r} - \frac{L^2}{r^3}. \tag{5.62}$$

If the rotation curve is relatively flat, then $F_{\rm eff}(r)$ evaluated at $r_a \simeq r_g + \Delta r/2$ is

$$F_{\rm eff} = \left.\frac{\mathrm{d}\Phi_{\rm eff}}{\mathrm{d}r}\right|_{r_a} = \frac{v_c^2}{r} - \frac{L^2}{(r_g+\frac{1}{2}\Delta r)^3} \simeq \frac{v_c^2}{r_g}\left(1-\frac{\Delta r}{r_g}\right) - \frac{L^2}{r_g^3}\left(1-\frac{3}{2}\frac{\Delta r}{r_g}\right) \tag{5.63}$$

$$= -\Delta r\left(\frac{v_c^2}{r_g^2} + \frac{3L^2}{2r_g^4}\right) = -\frac{5v_c^4\Delta r}{2L^2}, \tag{5.64}$$

and similarly when evaluated at r_p, but with opposite sign. Equation (5.60) then becomes

$$J_r \simeq \frac{1}{8}\sqrt{\frac{5}{2}}\frac{v_c^2\Delta r^2}{L}. \tag{5.65}$$

Differentiating the above expression, we find

$$\frac{\mathrm{d}J_r}{\mathrm{d}\Delta r} = \frac{1}{4}\sqrt{\frac{5}{2}}\frac{v_c^2\Delta r}{L} \tag{5.66}$$

$$= \sqrt{\frac{1}{2}\sqrt{\frac{5}{2}}\frac{J_r v_c^3}{L}}. \tag{5.67}$$

Hence, for a small change δJ_r we can estimate the corresponding change in the radial amplitude, $\delta\Delta r$, which is likely to be a good estimate for the positional error we would make in assuming that a star with action $J_r + \delta J_r$ actually had action J_r.

We can confirm the predictions of the above equations numerically. The left panel of Fig. 5.14 shows the real-space trajectory of the circular orbit I2, with radius $r = 8\,\mathrm{kpc}$, in the isochrone potential of

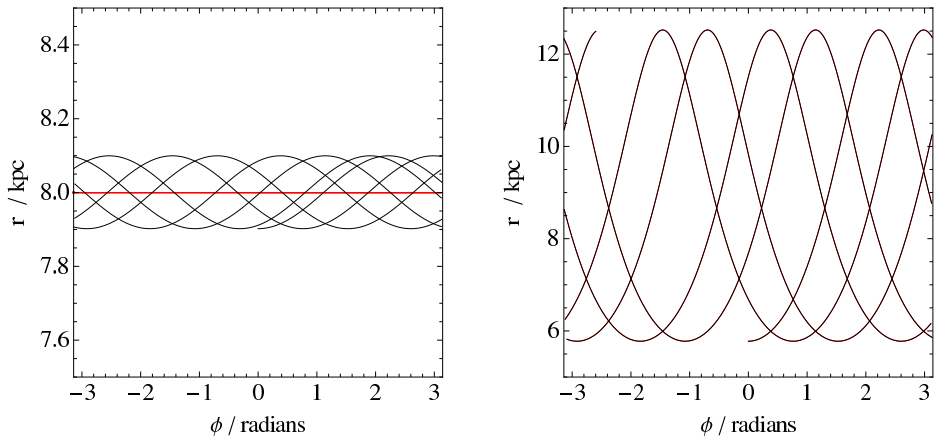

Figure 5.14: Plots to demonstrate the effect of changes in action J_r on the real-space orbital trajectory of streams. The left panel shows that a small change in J_r ($\delta J_r = 0.21\,\mathrm{km\,s^{-1}\,kpc}$) to a circular orbit (I2 in Table 5.2) produces a change of $\delta\Delta r \sim 0.2\,\mathrm{kpc}$ in radial amplitude. The right panel shows the effect of the same perturbation δJ_r on an eccentric orbit (I3 in Table 5.2) with the same angular momentum as the circular orbit, but with $J_r = 207\,\mathrm{km\,s^{-1}\,kpc}$. The effect on radial amplitude is so small as to be invisible in this plot.

Table 5.1. Also plotted is the trajectory of an orbit that has identical L to I2, but $J_r = 0.21\,\mathrm{km\,s^{-1}\,kpc}$.

Clearly, equation (5.67) ceases to have meaning when faced with orbits very close to circular, so we rely on the integral form given by equation (5.65) instead for our estimate. Equation (5.65) predicts $\Delta r = 0.19\,\mathrm{kpc}$ for this perturbation from circular, which appears from the left panel of Fig. 5.14 to be close to exact.

From equation (5.67) we see that the magnitude of the effect diminishes as $\delta\Delta r \sim 1/\Delta r \sim 1/\sqrt{J_r}$. The right panel of Fig. 5.14 shows this to be the case. The panel shows two trajectories in the isochrone potential: one for the orbit I3, and one for the same orbit with J_r incremented by 0.1 per cent. Equation (5.67) predicts a change $\delta\Delta R \sim 8\,\mathrm{pc}$. Close inspection of the trajectories confirms an actual $\delta\Delta r \sim 7\,\mathrm{pc}$, so the prediction is correct, but as is clear from Fig. 5.14, corrections of such magnitude are negligible.

What follows is an estimate for the positional error made by incorrectly guessing L. Consider again a cluster on an orbit close to circular. The angular momentum of the cluster is related to the guiding centre radius r_g and the circular velocity v_c by

$$L = v_c r_g. \tag{5.68}$$

Consider now a star whose angular momentum is suddenly reduced by δL. This star is now at apocentre,

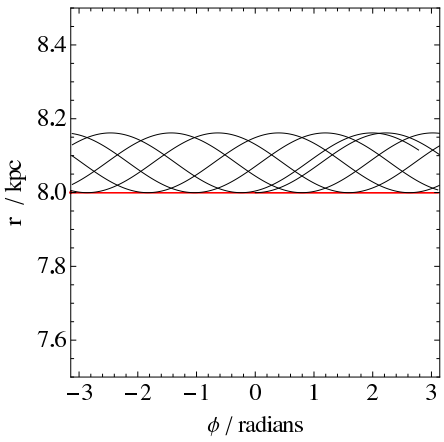

Figure 5.15: Plot to demonstrate the effect of changes in angular momentum L on real-space orbital trajectory. The trajectory of two orbits is plotted: the red line shows the trajectory of the orbit I2 (Table 5.2), while the black line shows the same orbit, except with L incremented by 1 per cent.

since its guiding centre radius has been reduced by

$$\delta r_g = \frac{\delta L}{v_c}, \tag{5.69}$$

where we have assumed that the rotation curve is flat. Pericentre radius will have been reduced by of order twice the change in guiding centre radius, hence we may write

$$\delta \Delta r = \frac{2\delta L}{v_c}. \tag{5.70}$$

Thus, we can predict the change in radial amplitude $\delta\Delta r$ for a small change in angular momentum δL, which is likely to be a good estimate for the positional discrepancy we would encounter in assuming that a star with angular momentum $L + \delta L$ actually had angular momentum L.

Again, we can check the predictions of this expression numerically. Fig. 5.15 shows the trajectories of two orbits in our isochrone potential. The red line is orbit I2, while the black line is the same orbit, but with the angular momentum increased by 1 per cent. Equation (5.70) predicts $\Delta r \simeq 0.16\,\mathrm{kpc}$ for this change. Fig. 5.15 shows that this estimate is close to exact.

The only example we have shown thus far where this effect could be of consequence is the isochrone-potential stream shown in Fig. 5.7. For this cluster, equation (5.67) predicts that a positional error of $\sim 3\,\mathrm{pc}$ would be accrued by assuming all stars have the same radial action. Similarly, equation (5.70) predicts that a positional error of $\sim 1.5\,\mathrm{pc}$ will be accrued by assuming that all stars have the same angular momentum. These errors are insignificant, so no corrections are required in this case.

However, we shall see in some later examples that the errors will not be negligible. In these cases, the

correction described by equation (5.53), properly accounting for variation in action down the stream, will be required. In such cases, we will assume that we accurately know t_d, the time since the first pericentre passage. In general we would not know t_d accurately, although given $\mathbf{J}_0$ we could make a reasonable guess as to its value. However, even a poor, but finite, guess for the value of t_d would likely produce a more accurate real-space stream track than would assuming $\mathbf{J} = \mathbf{J}_0$ everywhere along the stream.

5.5 The consequences of fitting orbits

We confirmed in §5.3.4 that streams formed in the isochrone potential do not necessarily delineate orbits. In this section, we briefly examine the consequences should one attempt to constrain the parameters of the potential by assuming that these streams *do* delineate orbits.

We construct an experiment as follows. We first create three sets of pseudo-data, each containing a set of phase-space coordinates corresponding to one of the three stream tracks in Fig. 5.10. We choose to consider the streams from Fig. 5.10 because they are at apocentre, and many actual observed streams (e.g. the Orphan stream of Belokurov et al. (2007)) are discovered close to apocentre.

We now wish to measure the quality with which an orbit for a given set of isochrone potential parameters can be made to fit the data. Unlike in Chapters 2 and 3, where orbits were reconstructed from data for which only partial phase-space information is available, for this exercise we have granted ourselves pseudo-data with full and accurate positional and velocity information. This simplifies considerably the matter of finding an orbit that is close to the best fitting one.

We choose an orbit as follows. We first select a datum near the centroid of the stream, and declare that our chosen orbit must pass directly through this datum. Although it may be that some nearby orbit, one that does not pass directly through this point, would make a better-fitting orbit, any such orbit must pass very close to the selected datum, because it is close to the centroid. Thus, such an orbit cannot be much better-fitting than one that passes directly through the datum. Having chosen a datum, for a given set of potential parameters, an orbit is defined.

Having chosen our orbit, a goodness-of-fit statistic χ^2 is calculated as follows. For each datum in the stream, with phase-space coordinate $\mathbf{w}_i$, a location along the orbit $\mathbf{w}'_i$ is chosen that minimizes the square difference

$$(\mathbf{w}_i - \mathbf{w}'_i)^2. \tag{5.71}$$

Having obtained the $\mathbf{w}'_i$, the goodness-of-fit χ^2 is defined by

$$\chi^2 = \sum_{i,j} \frac{(w_{i,j} - w'_{i,j})^2}{\sigma_j^2}, \tag{5.72}$$

where j are the phase-space coordinates, and σ_j is the rms of $(w_{i,j} + w'_{i,j})/2$ over i. This χ^2 statistic provides a dimensionless measure of the phase-space distance between the best-fitting orbit in a given potential, and the pseudo-data.

If the pseudo-data set were a sample of a perfect orbit in some potential, we expect the value of χ^2 to be exactly zero, when the correct potential parameters are considered. As the potential parameters are varied away from their true values, we expect the value of χ^2 to rise, as the best-fitting orbit becomes a steadily worse representation of the data. Hence, we expect minima in χ^2 to be associated with the potential parameters that are optimum, from the perspective of fitting an orbit to the data. We seek such minima by plotting contours of χ^2 over a range of likely values for the potential parameters.

5.5.1 Results

Three clusters of 50 test particles were created, with small initial angle-space distributions. The action-space distribution of each was a segment of a line, $\Delta J = 0.2\,\mathrm{km\,s^{-1}}\,\mathrm{kpc}$ long, and oriented such that, after a long time, the angle-space distribution of each would precisely match one of the three streams shown in Fig. 5.10.

Each cluster was placed on the orbit I1 in the isochrone potential of Table 5.1, and integrated forward for 77.04 Gyr to produce a thin stream about 10 kpc in length, that has its centroid at apocentre. The streams are then similar to those shown in Fig. 5.10. We declare each one of these streams to be a pseudo-data set for our experiment.

Fig. 5.16 shows contours for the goodness-of-fit χ^2, for the best-fitting orbit, in an isochrone potential with parameters GM and b as shown. In each case, a valley of low values, corresponding to a family of approximately degenerate solutions, can be seen. This degeneracy occurs because the trajectories (being at apocentre) are relatively insensitive to the shape of the rotation curve, but are very sensitive to the magnitude of the force. Hence, the family of solutions comprises those combinations of the (GM, b) parameters that give rise to the same circular velocity at the radius of the stream. In each plot, a red cross marks the correct parameters for the potential.

In the case of the low-curvature stream of the top panel of Fig. 5.16, for all values of b, the orbit fitting technique reports a lack of interior mass when compared with the middle panel. In the middle panel, the valley of solutions passes directly through the correct parameter coordinate, as expected. Thus, the performance of the orbit fitting technique is validated. In the case of the high-curvature stream of the

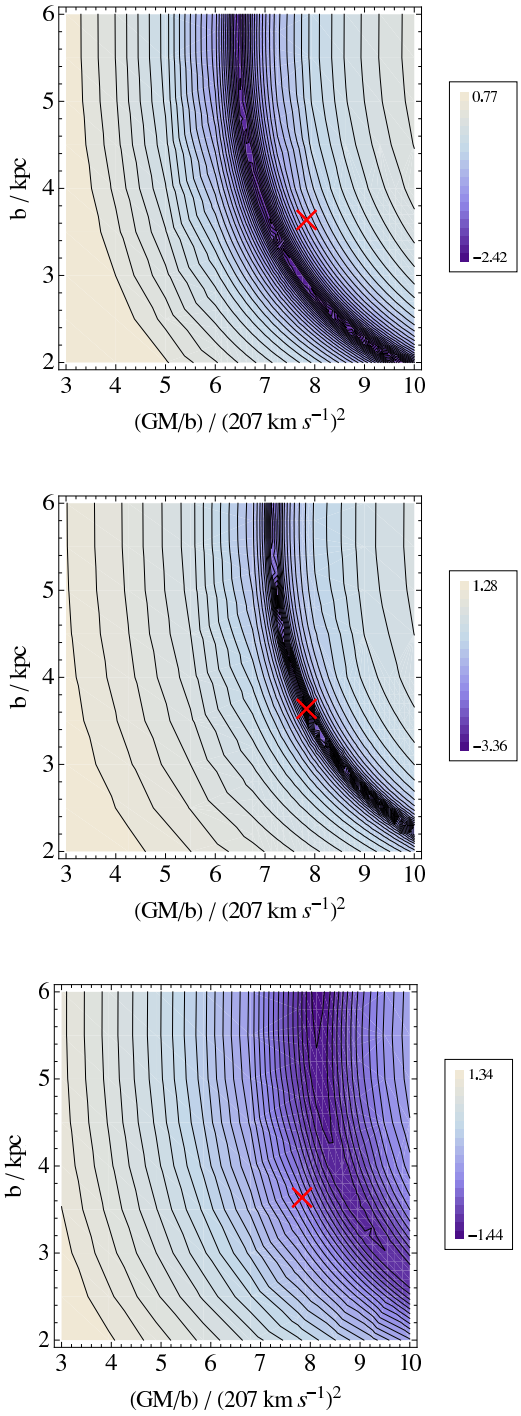

Figure 5.16: Plots of goodness-of-fit χ^2 for the pseudo-data to the best-fitting orbit in the potential, for a range of potential parameters $(GM/b, b)$. The pseudo-data sets are 50 particle realisations of those streams shown in Fig. 5.10. Top panel: the red stream from Fig. 5.10; middle panel: the black stream from the same figure; bottom panel: the blue stream from the same figure. In each case, a valley of low values, corresponding to approximately degenerate solutions, can be seen. The red cross marks the correct parameters for the potential. In the case of the low-curvature stream of the top panel, the orbit fitting technique reports a lack of interior mass when compared with the middle panel. In the middle panel, the valley of solutions passes directly through the correct parameter coordinate, as expected. In the case of the high-curvature stream of the bottom panel, the orbit-fitting algorithm reports an excess of mass interior to the stream when compared with the middle panel.

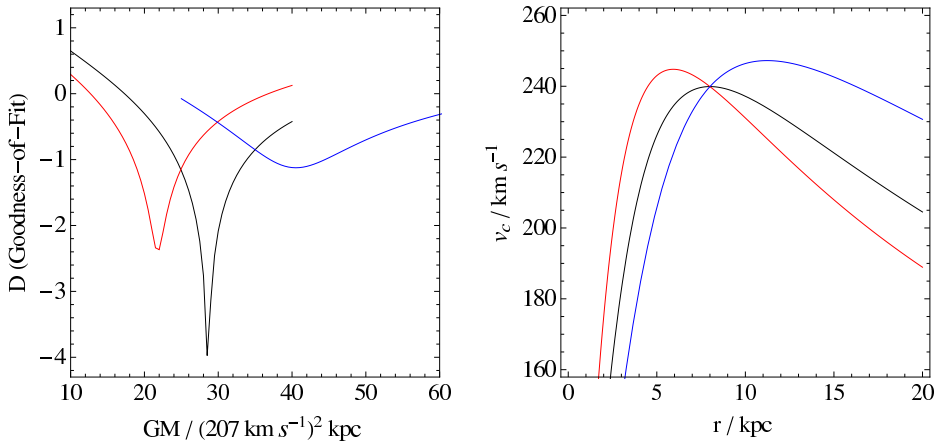

Figure 5.17: Left panel: goodness-of-fit χ^2 for each pseudo-data set, to the best-fitting orbit in the potential. The colours of the curves for each pseudo-data set identify it with the corresponding stream from Fig. 5.10 from which it was derived. Right panel: rotation curves for the optimum potential found for each pseudo-data set. Unlike in Fig. 5.16, we have required that the circular velocity $v_c = 240\,\mathrm{km\,s^{-1}}$ at $R_0 = 8\,\mathrm{kpc}$. This reduces the potential to one of a single parameter; in this case, mass. The quality of the fits to the red and blue data sets is significantly degraded. Similarly to Fig. 5.16, utilizing the low-curvature, red data set as a proxy for an orbit causes us to underestimate the host mass by approximately 21 per cent. Utilizing the high-curvature, blue data set causes us to overestimate the host mass by approximately 54 per cent.

bottom panel, the orbit-fitting procedure reports an excess of mass interior to the stream for all values of b, when compared with the middle panel.

Thus, the attempt to constrain the potential parameters by using misaligned streams as proxies for orbits has led us to err. In the case of the low-curvature stream, for any given value of b, we have underestimated the mass by about 11 per cent. In the case of the high-curvature stream we have overestimated the mass by approximately 14 per cent.

We do note that the lowest value of χ^2 seen for the two misaligned streams is significantly higher that the lowest value of χ^2 seen for the perfectly aligned stream. This indicates that in the case of the misaligned streams, although optimum values for the parameters (GM, b) are being found, the fit to the orbit is still not perfect there. This is precisely the effect utilized in Chapter 2 to try to identify the correct potential, and it is discussed there in further detail.

In the cases presented in Fig. 5.16, we have placed no constraints on the parameters of the potential other than those implied by the stream. This is a somewhat unrealistic test, since in practical usage one would generally require any acceptable potential to reproduce other observed features of the Milky Way galaxy, such as the circular velocity at the Solar radius. Most of the combinations of (GM, b) in the family of solutions do not reproduce the correct circular velocity. We therefore repeat the test, while considering only those combinations of parameters that do.

Fig. 5.17 shows the value of χ^2 obtained for the best fitting orbit, versus galaxy mass parameter

GM, while requiring b to take that value which gives the fiducial $v_c = 240\,\mathrm{km\,s^{-1}}$ at $R_0 = 8\,\mathrm{kpc}$. Like with Fig. 5.16, the quality of the fit at optimum GM is significantly degraded for the misaligned streams when compared to the perfectly aligned stream. The correct value of GM is obtained when fitting to the stream that perfectly delineates its orbit, thus validating the technique. However, the error in GM when deduced using the low-curvature stream is 21 per cent, and the error in GM when deduced using the high-curvature stream is 54 per cent. Hence, our attempt to use misaligned stream tracks to constrain the potential, while simultaneously requiring the potential to be consistent with other observations, has led us to yet greater error.

In conclusion, we find that there is a risk of substantial systematic errors in parameter estimation being made, if one attempts to constrain the potential using streams, and one assumes that streams perfectly delineate orbits.

5.6 The action-space distribution of disrupted clusters

Up until now, we have relied upon the qualitative estimate from §5.4.1 for what the action-space distribution of a disrupted cluster might be. In this section, we investigate the action-space distribution of various cluster models using N-body simulation. We further utilize our N-body models to confirm the misalignment between streams and orbits, and to demonstrate that we can accurately predict the real-space track of the stream.

The action-angle coordinates, as defined in the host galaxy potential, have limited usefulness when applied to the particles in a bound cluster, because the actions are not constant with time. Nonetheless, they remain a valid set of canonical coordinates and can be legitimately used to describe the phase-space distribution of the cluster. Moreover, we shall see that a disrupted cluster gives rise to a characteristic distribution in action-space, from which a precise track of the stream can be predicted.

Fig. 5.18 shows a segment of each of the orbits I4 and I5 in the isochrone potential of Table 5.1. These orbits were chosen to be fairly representative of those occupied by tidal streams in our Galaxy: they have apocentre radius $\sim 20\,\mathrm{kpc}$ and are moderately eccentric to allow for efficient tidal stripping. For our investigation, we wish to launch model clusters on each of these orbits, where the otherwise-stable cluster has been chosen such that its outermost stars will be torn away by tidal forces close to pericentre. The process by which we choose our model clusters is detailed in the next section.

5.6.1 Cluster models

We choose to work with King models (King, 1966; Binney & Tremaine, 2008) for our clusters, since these simple models are both easy to generate and are fairly representative of some observed globular

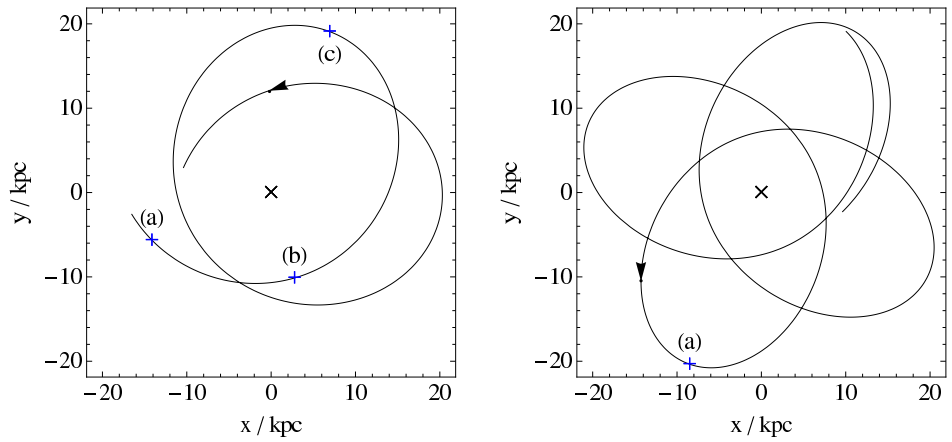

Figure 5.18: Plan views of the orbits used in this section. The left panel shows I4, with (a), (b) and (c) marking the positions of the cluster corresponding to the top-left, top-right and middle-left panels of Fig. 5.19 respectively. The right panel shows I5, with (a) marking the position of the cluster corresponding to the bottom-right panel of Fig. 5.23. In both panels, the potential in use was the isochrone potential described in Table 5.1.

cluster profiles (e.g. Fig. 6.18, Binney & Merrifield, 1998). The profile of a King model can be defined[4] by $W \equiv \Phi_0/\sigma^2$, the ratio of the central potential to the squared-velocity parameter σ^2. The resulting family of models have similar profiles in terms of $\rho(\tilde{r}/\tilde{r}_0)/\rho_0$, where ρ_0 is the central density and $\tilde{r}_0$ is the core radius. An exact model is specified by choosing ρ_0 and $\tilde{r}_0$, in addition to W, either directly or through a relation with another parameter.

For a given orbit, we specify our models as follows. Following the argument of Dehnen et al. (2004), we note that a cluster of mass M_c orbiting at a galactocentric radius r from the centre of a host galaxy with circular velocity v_c, will be tidally pruned to the cluster radius $\tilde{r}_{\rm tide}$, where

$$\tilde{r}_{\rm tide}^3 \simeq \frac{GM_c}{v_c^2} r^2. \tag{5.73}$$

We freely choose a profile parameter W and a cluster mass M_c, and we also specify a galactocentric stripping radius $r_s > r_p$, where r_p is the pericentre radius of the orbit concerned. We then set $\tilde{r}_t$, the cluster truncation radius, equal to $\tilde{r}_{\rm tide}$ from equation (5.73), where $r \to r_s$ and $v_c \to v_c(r_s)$. The resulting cluster will remain intact while $r \gg r_s$, but will have its outermost stars tidally stripped when $r \sim r_p$.

[4]Equivalently, one may specify the concentration parameter, $c \equiv \tilde{r}_t/\tilde{r}_0$, the ratio of the truncation radius $\tilde{r}_t$ to the core radius $\tilde{r}_0$. There is a one-to-one correspondence between c and W, though the form of the relation is not trivial. A plot of c versus W can be seen in Binney & Tremaine (2008, §4.3.3c)

Table 5.4: Details of the cluster models used in this section. Defining parameters are in the central block of columns, while derived parameters occupy the right-hand block.

	W	M_c	r_s	$r'_{\rm lim}$	σ	$t_{\rm dyn}$	ϵ
C1	2	$10^4 M_\odot$	12 kpc	48.6 pc	$1.18\,{\rm km\,s^{-1}}$	12.6 Myr	1.0 pc
C2	2	$10^5 M_\odot$	12 kpc	104.8 pc	$2.54\,{\rm km\,s^{-1}}$	12.6 Myr	2.2 pc
C3	6	$10^4 M_\odot$	12 kpc	48.6 pc	$1.14\,{\rm km\,s^{-1}}$	2.36 Myr	0.32 pc
C4	2	$10^4 M_\odot$	11 kpc	45.5 pc	$1.22\,{\rm km\,s^{-1}}$	11.79 Myr	0.94 pc
C5	2	$10^4 M_\odot$	11 kpc	45.7 pc	$1.22\,{\rm km\,s^{-1}}$	11.79 Myr	0.94 pc
C6	2	$10^4 M_\odot$	15 kpc	56.3 pc	$1.10\,{\rm km\,s^{-1}}$	15.6 Myr	1.2 pc
C7	2	$10^4 M_\odot$	12 kpc	48.2 pc	$1.19\,{\rm km\,s^{-1}}$	12.3 Myr	0.99 pc

5.6.2 The disruption of a cluster

The low-mass cluster model C1 (Table 5.4) was specified for the orbit I4 (Table 5.2) according to the schema in §5.6.1. We chose a low value for the profile parameter of $W = 2$ for our basic cluster model, in order to ensure the presence of many particles near the cluster truncation radius $\tilde{r}_t$ during successive stripping events.

A 10^4 particle realization of the cluster model C1 was made by random sampling of the King model distribution function (Binney & Tremaine, 2008, equation 4.110). This cluster was placed at a point shortly after apocentre on the orbit I4 in the isochrone potential of Table 5.1. The cluster was evolved forward in time in the aforementioned potential by the FVFPS tree code of Londrillo et al. (2003), using a time step of $\mathrm{d}t = t_{\rm dyn}/100$ and a softening length ϵ as specified in Table 5.4. The simulated time period was 4.81 Gyr, or almost 14 complete radial orbits.

Fig. 5.19 shows the evolution of the action-space distribution of the cluster model C1 as a function of time. In all the panels of that figure, an arrowed black line shows the mapping of the frequency vector from angle space into action space, $\mathbf{D}^{-1}\boldsymbol{\Omega}_0$. This vector shows the direction that maps onto $\boldsymbol{\Omega}_0$ in angle-space, so any action-space distribution that is aligned with this vector will be aligned with $\boldsymbol{\Omega}_0$ in angle-space.

In all cases involving the isochrone potential, $\mathbf{D}^{-1}\boldsymbol{\Omega}_0$ is oriented exactly along the J_r axis. We can understand this from equation (5.43), which shows that in the isochrone potential, the frequency direction $\hat{\boldsymbol{\Omega}}_0$ is a function of L only, and is independent of J_r. Thus, a line of constant $\hat{\boldsymbol{\Omega}}_0$ must map into action-space as a line of constant L. We note that this is a peculiar feature of the isochrone potential, and is not true for a general potential.

The upper-left panel of Fig. 5.19 shows the configuration of the cluster immediately after release, at position (a) in Fig. 5.18. The distribution is ellipsoidal but without additional substructure, which we expect since the distribution in action results entirely from the approximately spheroidal density profile, and the approximately isotropic velocity dispersion, of the cluster. We understand the ellipticity of the action-space distribution from §5.4.1. Indeed, the prediction of equation (5.52) of $\Delta J_r/\Delta L \sim 0.3$ for this

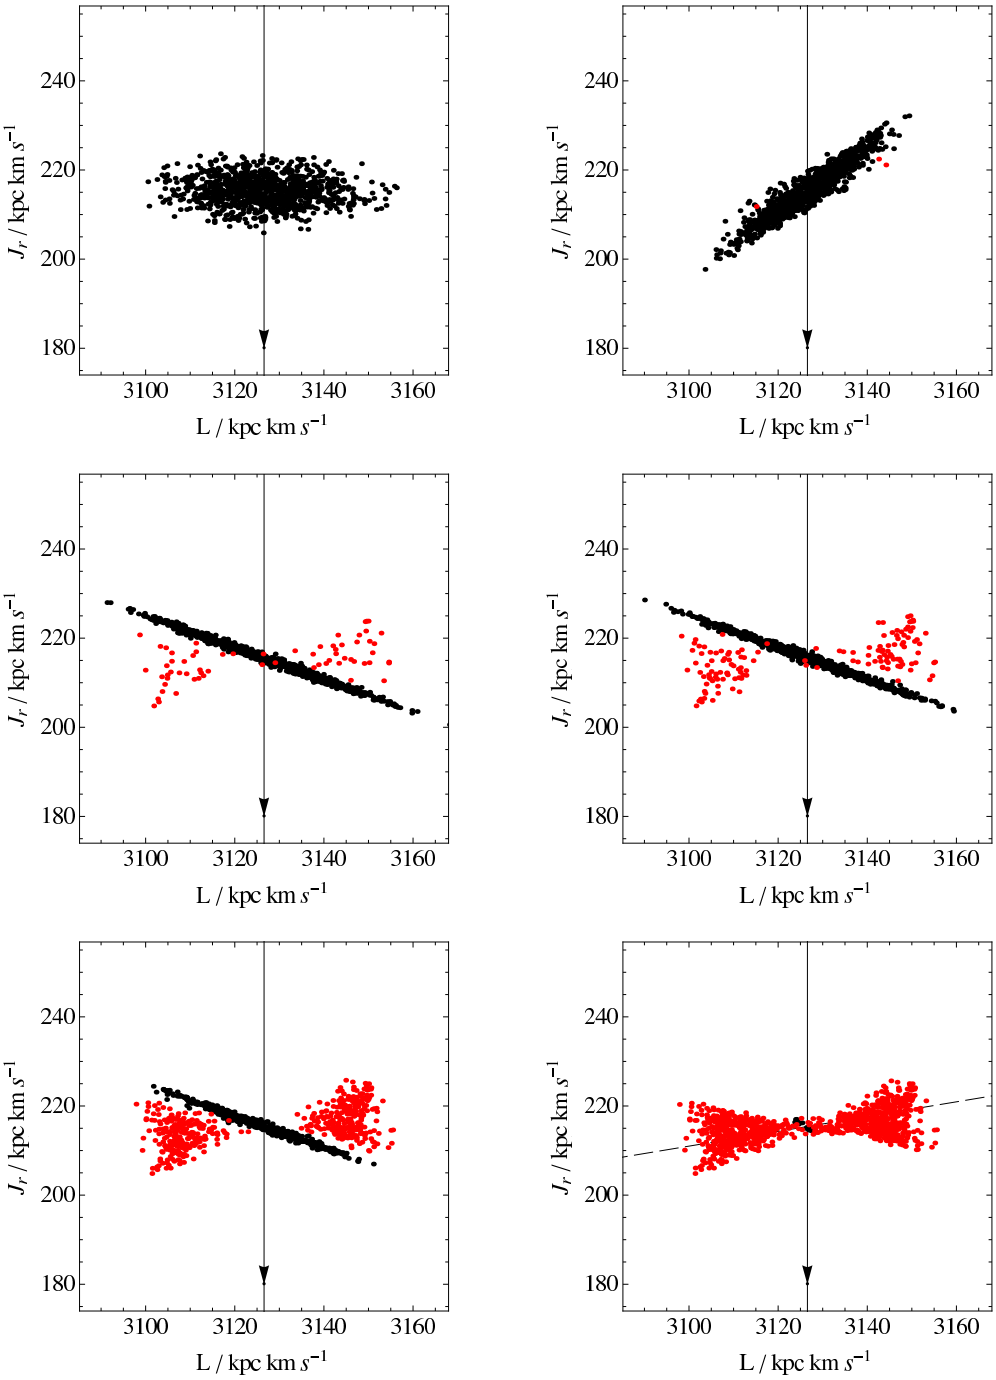

Figure 5.19: Plots showing the action-space distribution of particles for the N-body cluster model C1, at different times along the orbit I4. From left-to-right and top-to-bottom, these times are: shortly after release; first pericentre passage; first apocentre passage; second apocentre passage; 7th apocentre passage; 14th apocentre passage. The solid black line is the inverse map of the frequency vector, $\mathbf{D}^{-1}\boldsymbol{\Omega}_0$. The dashed line in the bottom-right panel represents a least-squares linear fit to the particle distribution.

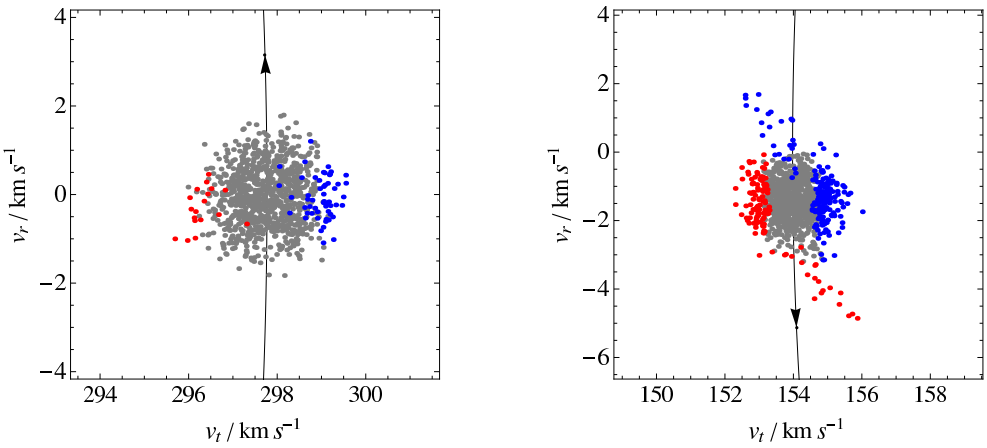

Figure 5.20: Plots of radial velocity against tangential velocity for the C1 cluster model, on the I4 orbit, at (left panel) the first pericentre passage, and (right panel) the first apocentre passage.

orbit can be seen to be approximately correct.

In the top-right panel of Fig. 5.19, the cluster has now moved from its point of release to its first pericentre passage, marked as position (b) in Fig. 5.18. We see that the ellipse has flattened somewhat, and has rotated anticlockwise. In the middle-left panel, the cluster has now progressed to the subsequent apocentre passage, marked as position (c) in Fig. 5.18. Here, the cluster is again flattened, but now it has rotated clockwise.

We can qualitatively understand this behaviour as follows. Consider a particle in a cluster near apsis. What changes to the actions will be made by perturbations to the velocity of this particle? A perturbation δv to the transverse velocity will cause a change to the angular momentum

$$\delta L = r\delta v. \tag{5.74}$$

By means of the mechanism described by equation (5.70), this δL will cause a change in the guiding centre radius r_g, which will cause a corresponding change in the radial action, according to equation (5.67). Conversely, a perturbation to the radial velocity will cause negligible change to the radial action, since

$$\delta E \simeq p_r \delta p_r = \dot{r}\delta v \sim 0, \tag{5.75}$$

and $J_r(E, L)$ remains unchanged. Hence, the distribution in both J_r and L is governed primarily by the transverse velocity, when the cluster is at apsis, and their highly correlated distribution reflects this.

We can confirm this analysis by examining Fig. 5.20. The left panel shows (v_r, v_t) for the cluster near its first pericentre passage, while the right panel shows the same for the cluster near its subsequent apocentre passage. The particles coloured red are those with $L < 3110\,\mathrm{kpc\,km\,s^{-1}}$ and those coloured

blue have $L > 3140\,\mathrm{kpc\,km\,s^{-1}}$. The boundary between the colours is very sharp in the v_t direction, as one would expect, since the particles have been coloured on the basis of L. However, we note that particles with the full range of v_r contribute to both the blue and red regions equally, despite Fig. 5.19 showing that J_r is very different for these two regions. Hence J_r must be independent of v_r near apsis.

We complete our explanation of the orientation of the action-space distribution by taking note of the sign of the change in J_r near apsis. Increasing L will always increase r_g, and when the cluster is at pericentre, this pushes r_g further away, and so increases the radial action. Conversely, when the cluster is at apocentre, increasing r_g brings it closer to the cluster, and so decreases the radial action. Thus, we see that near pericentre, particles with high L will have high J_r and the distribution will be rotated to have a positive gradient $\Delta J_r/\Delta L$. Conversely, at apocentre, particles with high L will have low J_r, and so the distribution will be rotated to have a negative gradient $\Delta J_r/\Delta L$.

We can predict the value of the gradient of this distribution, by combining equation (5.67) and equation (5.70). We find

$$\frac{\mathrm{d}J_r}{dL} \sim \pm\sqrt[4]{\frac{10J_r^2}{L^2}} \sim \pm\sqrt{\frac{\pi J_r}{L}}, \tag{5.76}$$

where the sign of the radical depends on the apsis under consideration, as detailed above. We can see from the upper-right panel of Fig. 5.19 that this equation predicts approximately the correct gradient, when evaluated for the orbit I4. Furthermore, we note that this equation implies that the gradient will be steeper for an orbit of greater eccentricity.

We now examine again Fig. 5.19, and note that the black particles are those still bound to the cluster, while the red particles are those that are unbound, where we have defined 'bound' to mean particles that are within r_{tide} of the cluster barycentre. In the middle-left plot, which shows the cluster configuration at the first apocentre passage, the few unbound particles form an approximately horizontal distribution. These are particles that have been stripped from the cluster near pericentre, i.e. in the top-right panel of Fig. 5.19.

We note two things. Firstly, the bulk of the stripped particles have large ΔJ from the cluster centroid; this is simply a consequence of the high-speed stars being most likely to be stripped. Secondly, we note that although the red particles in the middle-left panel span approximately the same range in L as the black particles in the top-right panel, they span a range in J_r that is only about half that spanned by the black particles.

We explain this as a result of the cluster's self-gravity, as follows. Particles that are stripped from the cluster mostly escape through the Lagrange points L_1 and L_2. As illustrated in Fig. 5.21, these two points are oriented along a radial that runs through the barycentres of the host galaxy and the cluster. Thus,

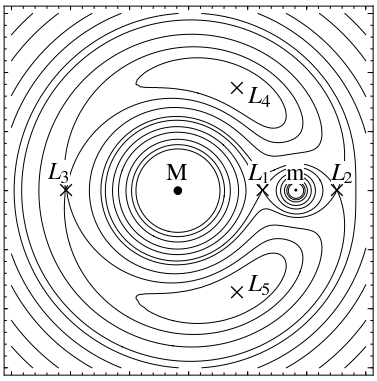

Figure 5.21: Lagrange points of neutral force, for the reduced three-body problem of a satellite of mass m on a circular orbit around a host galaxy of mass $M > m$. The lines are contours of constant effective potential $\Phi_{\rm eff}$, as given for this problem by equation (8.87) of Binney & Tremaine (2008), who also provide a diagram similar to this. In the example shown, the frame rotates with the orbiting masses and the entire configuration is static with respect to time. In the case of a cluster on an eccentric orbit, the Lagrange points will not be static, but the points L_1 and L_2 through which cluster particles escape are always aligned radially with respect to the cluster and the host galaxy.

it is the particles' radial velocity which initially carries them away from the cluster, and it is from this velocity component that the particles pay most of the energetic penalty for escaping, with the consequence that the radial velocity dispersion of the escaping stars is reduced. This reduction in the radial velocity dispersion corresponds to a compression of the distribution in J_r for escaping particles. Once unbound, the particles are carried further away from the cluster along a complex trajectory that ends up with the now-free particle drifting away from the cluster according to the mapping in equation (5.9). Thus, the final sum of the energetic penalty is paid from the difference in action $\mathbf{J} - \mathbf{J}_0$ between the particle and the cluster, with the net result that the unbound distribution uniformly shrinks. The latter effect is minor compared to the compression in J_r, because much more work is done in becoming unbound than in escaping to infinity once already unbound. Hence, the complete effect is to generate an unbound distribution, that looks like the high $\mathbf{\Delta J}$ wings of the pericentre distribution, but is compressed in J_r and shrunk slightly.

Looking again at Fig. 5.19, we note that in the bottom-left panel, which corresponds to the 7th apocentre passage, many particles have now escaped, and the size of the bound distribution has visibly shrunk. We understand this as a consequence of the most energetic particles having already escaped the cluster, leaving behind a colder core. By the time of the 14th apocentre passage, shown in the bottom-right panel, almost all the particles have escaped. We note that the positions of many of the red particles have remained static between the bottom-left and bottom-right panels. The action-space distribution of the unbound particles is therefore frozen in place, confirming that self-gravity is unimportant in streams.

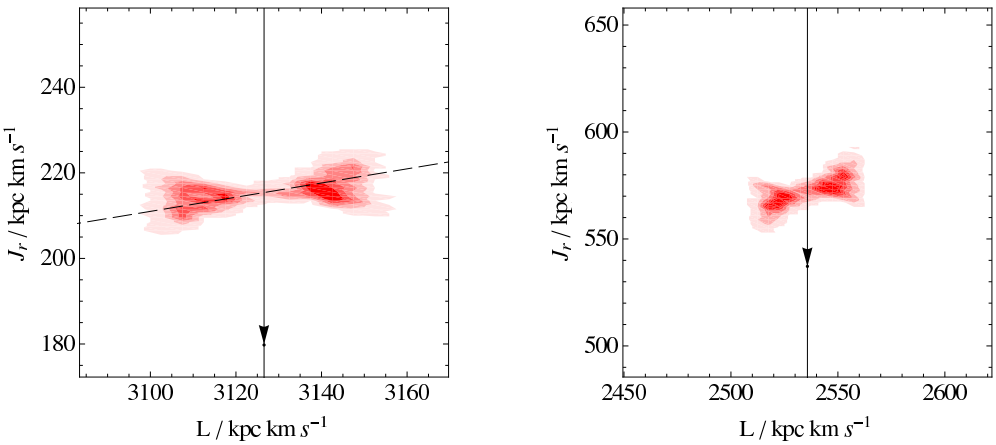

Figure 5.22: Plots showing the density of particles in action-space, corresponding to (left) the lower-right panel of Fig. 5.19, and (right) the lower-right panel of Fig. 5.23. The edges of the shaded areas represent contours of constant density, where the latter has been computed by placing particles into bins of width $\sim 2\,\mathrm{kpc\,km\,s^{-1}}$. Darker shading represents regions of higher particle density: the gross "bow-tie" structure is clearly seen, but no sub-structure is visible in either plot that is obviously unattributable to sampling noise.

The scatter plots in the bottom panels of Fig. 5.19 are filled too densely with particles to allow proper examination of the variation in particle density within the distribution. In order to elucidate this variation, the left panel of Fig. 5.22 shows a plot of particle density for the bottom-right panel of Fig. 5.19. The density field exhibits the same gross "bow-tie" structure that is visible in Fig. 5.19, and it further shows that particles are concentrated towards the centres of the lobes of the distribution, with particle density remaining low near the centroid. We explain this simply as a consequence of the number-density distribution of particles in velocity-space: the velocity-centroid of a cluster is populated with few particles because, although the phase-space density of particles there is high, the real-space volume associated with the lowest velocities is vanishingly small as one approaches the centroid, and so the velocity-space density is low there.[5] This aside, no significant fluctuations in particle density can be seen in the left panel of Fig. 5.22 that are obviously unattributable to the effects of sampling noise.

In conclusion, we have qualitatively understood the distortion of a cluster in action-space as it passes through pericentre and apocentre along its orbit. We have found that stars stripped at pericentre form a distribution that is derived from the pericentre distribution of bound stars, but is compressed in J_r. We have found that the pericentre distribution will exhibit a high correlation between J_r and L, and that the gradient of this correlation in (L, J_r) scales as $\sim \sqrt{J_r/L}$. We typically assume that our cluster orbit will have large L and comparatively smaller J_r: this would only be untrue for extreme plunging orbits which are not likely to be relevant to the problem in hand.[6] Hence the gradient in (L, J_r) will typically be less

[5] Precisely the same effect is seen in the Maxwell-Boltzmann distribution of particle speeds in a gas.

[6] A cluster on such an orbit would encounter severe difficulties while passing close the Galactic centre, where it would experience substantial tidal forces from the supermassive black hole. However, it is unlikely to even make it that far: since

than unity, and the compression will only act to shrink it still further. Hence, the stripping mechanism always results in an action-space distribution that is both flattened and very roughly oriented along $\hat{L}$.

5.6.2.1 The effect of changing the cluster model or orbit

We now investigate the qualitative effects of changing the cluster model parameters, or the cluster orbit, on the action-space distribution of a disrupted cluster. The cluster models used in this section are C1 to C4, detailed in Table 5.4. These clusters were created according to the schema of §5.6.1, taking the orbit to be I4 in the isochrone potential of Table 5.1 for the models C1–C3, and taking the orbit to be I5 in the same potential for the model C4.

The cluster model C1, the evolution of which along the orbit I4 was detailed in the previous section, is used as our baseline to which we compare the distributions of the other clusters. The cluster model C2 has the same profile parameter, $W = 2$, as does C1, but is 10 times more massive. The result is a cluster that is both heavier and proportionately larger while being stripped at the same galactocentric radius, r_s.

The cluster model C3 has the same mass as does C1, but is considerably more concentrated[7], with a profile parameter $W = 6$. The cluster has an identical truncation radius $\tilde{r}_t$ and velocity scale σ, but has significantly fewer particles near to $\tilde{r}_t$, when compared with C1. The particles of C3 are generally more tightly bound to the cluster than are those of C1.

The cluster model C4 has the same mass and profile parameter as C1, but is specified for the orbit I5, which has lower L than I4, and thus a smaller pericentre radius. The resulting cluster is slightly more compact, allowing it to survive to a closer galactocentric radius than can C1.

A 10^4 particle realization of each of the models C1–C3 was placed at a point shortly after apocentre on the orbit I4, and evolved forward in time by the FVFPS tree code, using a time step of $\mathrm{d}t = t_{\mathrm{dyn}}/100$ and a softening length ϵ as specified in Table 5.4. The total period of the simulation was 2.36 Gyr, or almost 7 complete radial orbits. Additionally, a 10^4 particle realization of C4 was placed at a point shortly after apocentre on the orbit I5, and evolved forward in time by the FVFPS tree code, for a total period of 2.21 Gyr, or almost 7 complete radial orbits.

Fig. 5.23 shows the action-space distribution for the cluster at the seventh apocentre passage, for each of these simulations. The top-left panel shows C1 on the orbit I4, and is identical to the bottom-left panel in Fig. 5.19. The top-right panel shows the cluster C2 on the same orbit; the bottom-left panel shows the cluster C3 on the same orbit; and the bottom-right panel shows the cluster C1 on the orbit I5.

it must feel the maximum possible tidal force from the Galaxy at some point along its plunging orbit, it is likely to break-up well away from pericentre.

[7]Fig. 4.8 of Binney & Tremaine (2008) shows the density profile for King models with a variety of values of W.

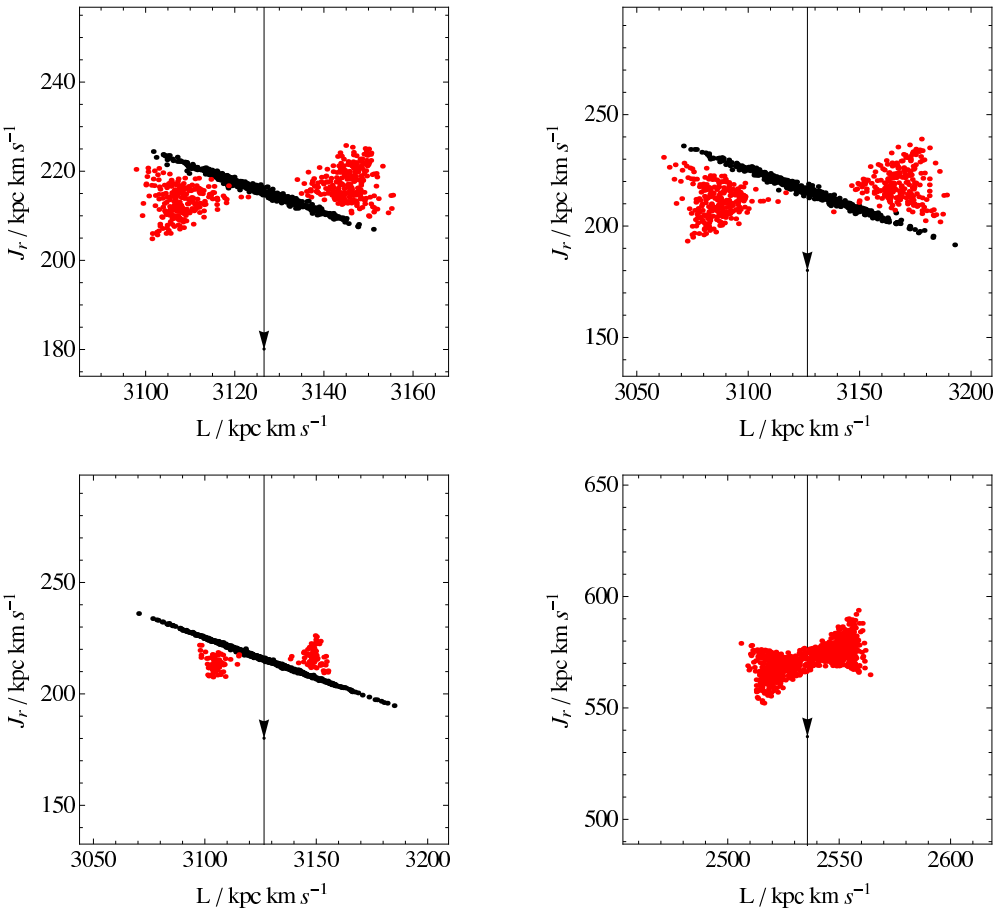

Figure 5.23: Plots showing the action-space distribution of particles for cluster models on various orbits. Each plot shows the distribution at the 7th apocentre passage. Black points are bound to the cluster, red points orbit free in the host potential. From left-to-right and top-to-bottom, the panels show: the C1 cluster on the orbit I4; the C2 cluster on the orbit I4; the C3 cluster on the orbit I4; and the C4 cluster on the orbit I5.

In the top-right panel, we see that the action-space distribution of the more massive cluster is qualitatively identical to that of C1, except that the distribution is approximately twice the scale. We conclude that cluster mass plays little role in determining the structure of the action-space distribution of disrupting clusters, but can determine the scale.

The bottom-left panel shows the distribution from the more centrally concentrated cluster C3. In this case, the shape of the distribution of unbound particles is approximately the same as for C1. However, the scale is much smaller: with more particles being near the core of the cluster, they have to work harder to escape, resulting in a colder action-space distribution. However, the distribution is still qualitatively similar to that of C1. Thus we conclude that cluster concentration can determine the scale of the action-space distribution, but not the general shape.

The bottom-right panel shows the distribution from the cluster C4 but on the orbit I5, which has the same apocentre radius as I4 but a pericentre radius about 33 per cent smaller. Unlike in the other panels, the cluster has become completely unbound by the 7th apocentre passage on I5, which is likely to be a result of $(r_s - r_p)$ being slightly larger for C4 on I5, when compared to the other clusters on I4, resulting in more efficient stripping at pericentre.

The distribution shown in this plot has approximately the same scale in ΔL as does the distribution from I4, but has approximately twice the scale in ΔJ_r. We can understand this, on account of the $\Delta J_r/\Delta L$ described by equation (5.76) being steeper for I5 than for I4. Further, since we have already noted that this cluster was stripped faster than was C1, the energetic penalty for escaping must be lower, and so we expect less compression in J_r. The resulting distribution is similar to that of the top-left panel, but less compressed in the $\hat{J}_r$ direction. Thus we conclude that changing the cluster orbit can distort the shape of the action-space distribution, but does not affect its basic structure.

The bottom-right panel is filled too densely with particles to gauge properly the density variation within it. We therefore provide a particle-density plot for the bottom-right panel, which is shown in the right panel of Fig. 5.22. This density plot correctly reproduces the gross structure observed in the scatter plot and reveals that the particles are concentrated in the lobes of the distribution, as was also true for the density plot in the left panel of Fig. 5.22. Aside from this, the plot reveals no further density fluctuations in the action-space distribution that cannot be attributed to sampling noise.

5.6.3 Predicting the stream from the action-space distribution

We have determined the action-space distribution for several disrupted cluster models by means of N-body simulation. We now ask whether we can accurately predict the real-space path of the stream, given the action-space distribution of one of those models.

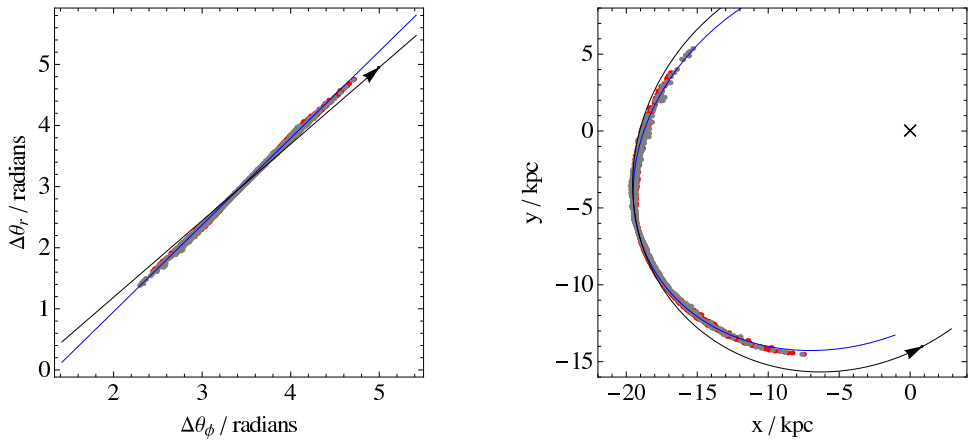

Figure 5.24: Plots of the distribution of particles for cluster C1, near its 14th apocentre passage on orbit I4. The left panel shows the angle-space distribution, while the right panel shows the configuration in real-space. The grey particles show positions directly computed from the N-body simulation, while the red particles show those positions predicted from mapping the action-space distribution in the bottom-right panel of Fig. 5.19. The two distributions almost precisely overlap. In both plots, the black arrowed curves show the trajectory of the progenitor orbit, while the blue curves show the mapping of the dashed line from Fig. 5.19. The blue curve is everywhere a much closer match to the stream particles than is the progenitor orbit.

Suppose that we know the time t_p since a cluster's first pericentre passage. The angle-space distribution is predicted by equation (5.9), where $\mathbf{D}$ is evaluated on the progenitor orbit $\mathbf{J}_0$ and $t \to t_p$. We will use as an example the action-space distribution shown in the bottom-right panel of Fig. 5.19, which corresponds to the 14th pericentre passage of the simulated cluster C1 on the orbit I4. The left panel of Fig. 5.24 shows the angle-space configuration corresponding to this panel: the grey particles are for angles directly computed from the results of the N-body simulation, while the red particles are for those predicted by equation (5.9), assuming t_p is known. Also plotted is the frequency vector $\mathbf{\Omega}_0$, shown as a black arrowed line. The distributions of black and grey particles in this panel agree perfectly. Furthermore, both distributions are obviously misaligned with the progenitor orbit.

The right panel of Fig. 5.24 shows the real-space configuration of particles for the same scenario. The grey particles are again plotted directly from the results of the N-body simulation, while the red particles result from the mapping into real-space of the red particles from the left panel. As in angle-space, the two distributions agree perfectly. Furthermore, the real-space manifestation of the misalignment of the stream with the orbit can be seen: the stream delineates a track that has substantially lower curvature than the orbit.

Our attempt to predict the real-space stream configuration from the action-space distribution has been completely successful. However, any complete model of the bottom-right panel of Fig. 5.19 must necessarily be rather complicated. It might not be possible to guess the form of this distribution without

full N-body modelling. The dashed line in the bottom-right panel of Fig. 5.19 is a least-squares fit of the action-space distribution to a line. This represents a rather more simple model of the action-space distribution, which it might well be possible to guess *ab initio* for a cluster on a given orbit.

How good a prediction for the stream track can we get from this line? The blue lines in Fig. 5.24 show the results of mapping this line into both angle-space and then real-space. It is clearly an excellent fit to the stream, in marked contrast to the orbit, which represents a very poor model of the stream by comparison. Thus, even a very simple model of the cluster in action-space—albeit one deduced from an accurate knowledge of the distribution—allows us to predict stream tracks accurately.

Finally, we note that during the mapping of this line into real-space, we need to make a correction to the cluster's action, as described by equation (5.53), to account for the variation in action down the stream. Evaluating equation (5.67) and equation (5.70) for the orbit I4, and taking $\mathrm{d}L \sim 25\,\mathrm{kpc\,km\,s^{-1}}$ and $\mathrm{d}J_r \sim 10\,\mathrm{kpc\,km\,s^{-1}}$ from the bottom-right panel of Fig. 5.19, we predict errors in the real-space of up to $\sim 0.2\,\mathrm{kpc}$ on account of the finite L distribution, and up to $\sim 0.15\,\mathrm{kpc}$ on account of the finite J_r distribution. These errors would be serious enough to be seen in Fig. 5.24, and thus the correction is required. We note, however, that even these substantial errors are insignificant compared to the several-kpc discrepancy between the stream and the orbit.

5.7 Non-spherical systems

We have investigated the formation of streams in spherical potentials, and demonstrated that they do not necessarily delineate streams, but that we can accurately predict their paths. Unfortunately, many real stellar systems in the Universe are not spherical. In particular, our own Galaxy, whose potential we are interested in probing with streams, is probably significantly flattened. In this section, we investigate the formation of streams in flattened potentials, and in particular we ask by how much they are misaligned with orbits. We also demonstrate that our apparatus is capable of correctly predicting stream tracks in flattened potentials, just as it is for spherical ones.

5.7.1 The general case with three actions

By analogy with §5.3.1, we now consider the case of a general axisymmetric stream-forming system, which is described by a Hamiltonian in three actions, $H(J_1, J_2, J_3)$. Equation (5.16) can be written

explicitly as

$$\begin{aligned}\Omega_1 e_{n,1} + \Omega_2 e_{n,2} + \Omega_3 e_{n,3} &= e_{n,1}\,\partial_1 H + e_{n,2}\,\partial_2 H + e_{n,3}\,\partial_3 H \\ &= \lambda(e_{n,1}\,\delta J_1 + e_{n,2}\,\delta J_2 + e_{n,3}\,\delta J_3) + k_n, \qquad (n = 1, 3). \end{aligned} \tag{5.77}$$

Defining the constants $\alpha_n = e_{n,2}/e_{n,1}$, $\epsilon_n = e_{n,3}/e_{n,1}$, $\beta_n = k_n/e_{n,1}$ we have

$$\partial_1 H + \alpha_n \partial_2 H + \epsilon_n\,\partial_3 H = \lambda_n(\delta J_1 + \alpha_n \delta J_2 + \epsilon_n\,\delta J_3) + \beta_n. \tag{5.78}$$

As in the spherical case, we can find a general solution to the homogeneous form of this equation by considering an arbitrary function f of the characteristic coordinates $(\delta J_1 - \delta J_2/\alpha_n)$ and $(\delta J_1 - \delta J_3/\epsilon_n)$. We add to this a particular solution for the inhomogeneous equation (5.78), to give the general solution

$$H = f\left(\delta J_1 - \frac{\delta J_2}{\alpha_n}, \delta J_1 - \frac{\delta J_3}{\epsilon_n}\right) + \beta_n \delta J_1 + \frac{\lambda_n}{2}\left(\delta J_1^2 + \delta J_2^2 + \delta J_3^2\right). \tag{5.79}$$

A system with a Hamiltonian of the form equation (5.79) is expected to exhibit globally consistent stream-forming geometry, described by the quantities $(\alpha, \beta, \lambda, \epsilon)_n$. Unfortunately, apart from trivial extensions of two-action systems such as the Kepler potential, there is no known Hamiltonian of the form equation (5.79) that is of relevance to galactic dynamics. Indeed, there is no known closed form for *any* non-trivial Hamiltonian in three actions that is relevant to galactic dynamics.

In principle, one would expand f as a power series in its two variables, and relate the resulting coefficients to a Taylor expansion of the Hamiltonian, as was done for the two-action case in §5.3.1. However, in the absence of a suitable analytic case to study, we will proceed no further. In the remainder of the section, the values of β_n and λ_n will be computed directly from $\mathbf{D}$ by numerical means, which itself will be constructed numerically from non-algebraic expressions for the frequencies and their derivatives.

5.7.2 Stäckel potentials

The only known integrable systems of three actions are those described by Stäckel potentials (Binney & Tremaine, 2008, §3.5.3). An exhaustive treatment in the context of galaxy modelling is given in de Zeeuw (1985) and de Zeeuw et al. (1986).

Stäckel potentials are of particular interest to us because, in the appropriate coordinate system, the Hamilton-Jacobi equation is separable. Action-angle variables can therefore be defined for these systems in terms of integrals over a finite path. Paul Stäckel first showed that the appropriate coordinate system is that of confocal ellipsoidal coordinates, and that indeed this is the only coordinate system in which the

Hamilton-Jacobi equation separates (de Zeeuw, 1985; Binney & Tremaine, 2008, p.228). It is in keeping with this that we note that Cartesian, spherical polar and cylindrical polar coordinates are themselves limiting cases of these coordinates.

We refer the reader to de Zeeuw (1985) for a detailed treatment of confocal ellipsoidal coordinates. We merely note here that we are adopting the de Zeeuw (1985) notation convention, and that we are restricting ourselves to considering only oblate axisymmetric potentials. In this case, the coordinate system becomes the prolate spheroidal coordinates (λ, ϕ, ν), where (λ, ν) are the roots for τ of

$$\frac{R^2}{\tau + \alpha} + \frac{z^2}{\tau + \gamma} = 1, \tag{5.80}$$

where (α, γ) are scaling constants, and $q = \sqrt{\gamma/\alpha}$ is the potential shape parameter, and (R, ϕ, z) are the familiar cylindrical polar coordinates. The permitted range for (λ, ν) is given by the inequality

$$-\gamma \leq \nu \leq -\alpha \leq \lambda. \tag{5.81}$$

In the meridonal plane (R, z), the coordinates (λ, ν) define a set of elliptic coordinates, while in the equatorial plane, (λ, ϕ) define a set of polar coordinates. A plot of the curves of constant (λ, ν) in the meridonal plane is shown in Figure 24 of de Zeeuw (1985). Of note is that curves of constant λ are a family of confocal ellipses, and curves of constant ν are a family of confocal hyperbolas, and that everywhere the two sets of curves are orthogonal.

In limit of $\lambda \to -\alpha$, we note that the direction $\hat{\lambda}$ of increasing λ is tangent with the radial vector $\hat{R}$ in cylindrical polar coordinates, while the direction $\hat{\nu}$ is tangent with the axial vector $\hat{z}$. We also note that in the limit of $\lambda \gg -\alpha$, the direction $\hat{\lambda}$ is tangent with the radial vector $\hat{r}$ in spherical polar coordinates, while the direction $\hat{\nu}$ is tangent with the polar vector $\hat{\theta}$.

The character of orbits in these potentials is as follows. The orbits circulate in ϕ. The orbits nutate in λ between two apses, where $-\alpha \leq \lambda_{\min} \leq \lambda \leq \lambda_{\max}$. The orbits bounce in ν, with a floor of $\nu_{\min} = -\gamma$ corresponding to $z = 0$, and an apex of $\nu = \nu_{\max}$ corresponding to $z = \pm z_{\max}$. Consecutive excursions in ν take place sequentially above and then below the equatorial plane: this is a consequence of the degenerate ellipsoidal coordinate system.

The limit of $\lambda \gg -\alpha$ will almost always apply for the example orbits that we examine below. We can qualitatively understand the meanings of the actions $(J_\lambda, J_\phi, J_\nu)$ in this limit as follows. J_λ is a generalization of the radial action J_r in spherical systems. Thus, an orbit with large J_λ will be eccentric, while an orbit with null J_λ will be confined to an ellipsoidal shell centred on the origin. An orbit with large J_ν will make excursions above and below the equatorial plane, while an orbit with null J_ν is

confined to the equatorial plane. Since we are only considering axisymmetric potentials, then $J_\phi = L_z$, the z-component of angular momentum, always. Our procedure for calculating these actions, and their corresponding angles, from conventional phase-space coordinates is discussed further in §5.7.4 below.

Finally, we mention the form of the potential. In prolate ellipsoidal coordinates, an oblate axisymmetric Stäckel potential takes the form (de Zeeuw et al., 1986),

$$\Phi(\lambda, \nu) = -\frac{(\lambda + \gamma)G(\lambda) - (\nu + \gamma)G(\nu)}{\lambda - \nu}, \tag{5.82}$$

where de Zeeuw's function $G(\tau)$ is determined once the density profile $\rho(z)$ along the z-axis has been chosen (see e.g. equation 23, de Zeeuw et al., 1986). Thus, the model is completely specified upon choosing $\rho(z)$ and the scaling parameters (α, γ).

5.7.3 Galaxy models with Stäckel potentials

de Zeeuw et al. (1986) shows that if one requires the density everywhere to be non-negative it is not possible to write down a Stäckel model in which the density $\rho(r)$ falls off with distance from he z-axis more rapidly than r^{-4} as $r \to \infty$. This is because an elementary density on the z-axis $\rho(z) = \delta(z - z_0)$ provides an off-axis density term that falls as r^{-4}. This behaviour rules out many classes of galaxy models, including disks with exponentially falling density profiles. However, we can construct models in which the density falls more slowly than r^{-4} as $r \to \infty$. In particular, models with asymptotically flat rotation curves, i.e. those in which $\rho(r) \sim r^{-2}$, are allowed.

In the models used in this section, we specify the z-axis density profile

$$\rho_z(z) = \frac{-\gamma\rho_0}{(z^2 - \gamma)} = \frac{-\gamma\rho_0}{\tau}, \tag{5.83}$$

where we have made use of $z^2 = \tau + \gamma$. In this case, de Zeeuw's function $G(\tau)$ can be written in closed form (equation 49, de Zeeuw et al., 1986). Models specified by equation (5.83) become spherical at large radii and have a rotation curve that is asymptotically flat, with

$$\lim_{r\to\infty} v_c^2 = -4\pi G\rho_0\gamma. \tag{5.84}$$

In the core of these models, the surfaces of constant density are approximately ellipsoidal, with axis

Table 5.5: Parameters for the example Stäckel potentials used in this chapter.

	$\rho_0/10^{10} M_\odot/\,\mathrm{kpc}^3$	$-\alpha/\,\mathrm{kpc}^2$	$-\gamma/\,\mathrm{kpc}^2$
SP1	0.361	29.64	8.89×10^{-3}
SP2	0.266	1.893	0.322

ratio[8]

$$\frac{a_z^2}{a_R^2} = \frac{2q^2}{(1-q^2)^2}\left(1 - q^2 + q^2 \log q^2\right), \tag{5.85}$$

where the central potential axis ratio $q = \sqrt{\gamma/\alpha}$. The models are completely specified by choosing a shape with q, a mass scale with ρ_0, and a distance scale with γ. The combination of an asymptotic logarithmic 'halo' and a flattened 'disk' in these models allow them to make a fair representation of the observed properties of disk galaxies, although the lack of freedom in the models severely restricts the shape of the flattened density profile that can be achieved.

Table 5.5 describes two such models. SP1 was chosen to simulate the highly flattened potential that may be felt in proximity to a heavy disk. The density axis ratio is fixed to be 10 near the solar radius of $R_0 = 8\,\mathrm{kpc}$, which is approximately the same ratio as for the (exponential) thin disk profile of the Milky Way (see Binney & Tremaine, 2008, Table 2.3). The specification is completed by requiring the rotation curve to peak at $R_0 = 8\,\mathrm{kpc}$ with a circular speed $v_c = 240\,\mathrm{km\,s^{-1}}$.

We note that the asymptotic circular velocity in this model is $v_c = 42\,\mathrm{km\,s^{-1}}$, which can be regarded as the halo contribution to the circular speed. The model is too centrally concentrated, and the halo contribution is too weak, to realistically model the Milky Way. However, it is highly flattened, and so makes an interesting example in which to study stream geometry. The rotation curve for this model is plotted as the blue curve in Fig. 5.1, while contours of constant density and potential for this model are shown in the left panels of Fig. 5.25.

The SP2 model was chosen to provide a force field better matched to that of the Milky Way at those radii where streams are typically observed, while still being somewhat flattened near the plane. The model was required to have $v_c = 240\,\mathrm{km\,s^{-1}}$ at $R_0 = 8\,\mathrm{kpc}$, and $v_c \geq 235\,\mathrm{km\,s^{-1}}$ at $R = 20\,\mathrm{kpc}$. The rotation curve was also required to peak at a radius not larger than $R = 5\,\mathrm{kpc}$. Finally, the model was required to be as flat as possible, subject to satisfying these constraints. These requirement completely specify the model, which has an asymptotic circular velocity of $v_c = 215\,\mathrm{km\,s^{-1}}$. The disk contribution to this model is therefore arguably rather weak, but it is in most respects a more reasonable model for the Milky Way galaxy than is SP1. The rotation curve for SP2 is plotted as the red curve in Fig. 5.1,

[8]Equation 46 in de Zeeuw et al. (1986) presents a formula for this quantity, but on inspection it must be wrong: it permits the density axis ratio a_z/a_R to take all values between $(0, \infty)$ for $q = (0, 1)$, and yet the formula is derived under the requirement that the central density profile is always oblate. The correct form for the expression is presented here, calculated directly from equation 35 in de Zeeuw et al. (1986).

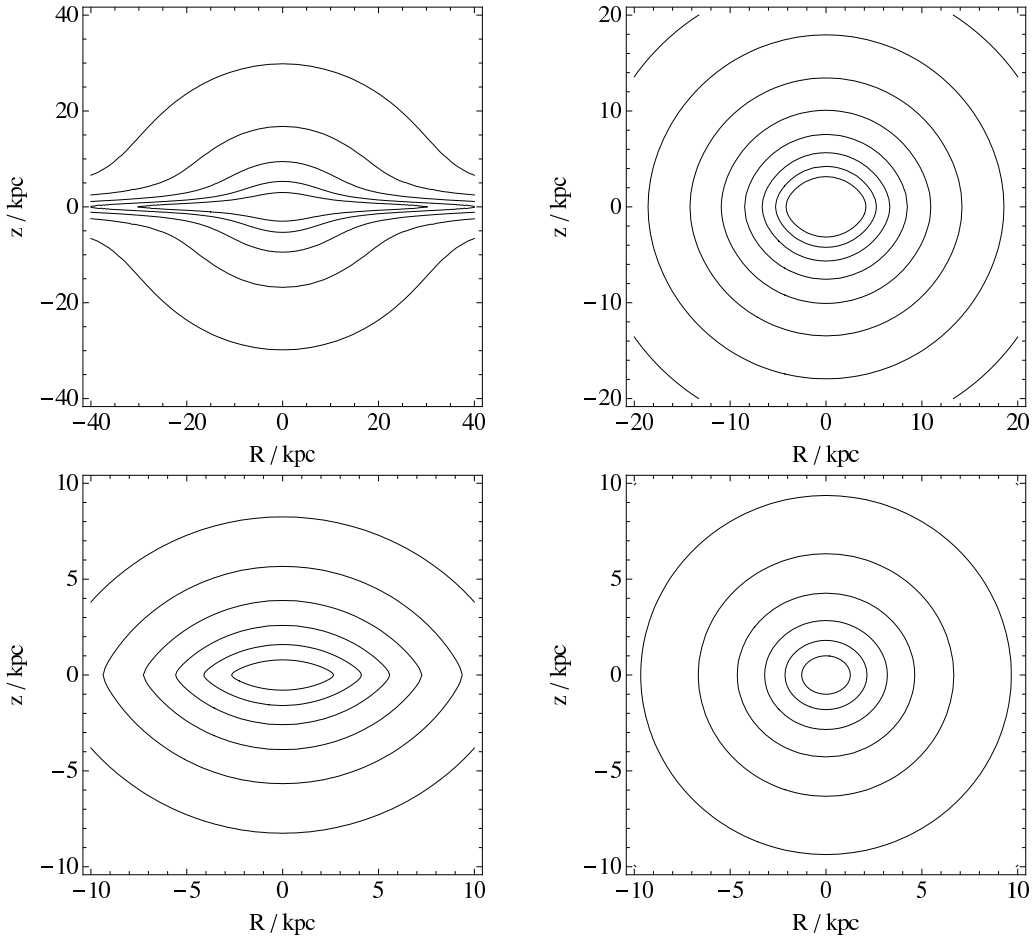

Figure 5.25: Stäckel models in use in this chapter. The left panels show the flattened model SP1, while the right panels show the rounder model SP2. The top panels show contours of $\log \rho/\rho_0$, while the bottom panels show contours of the potential Φ.

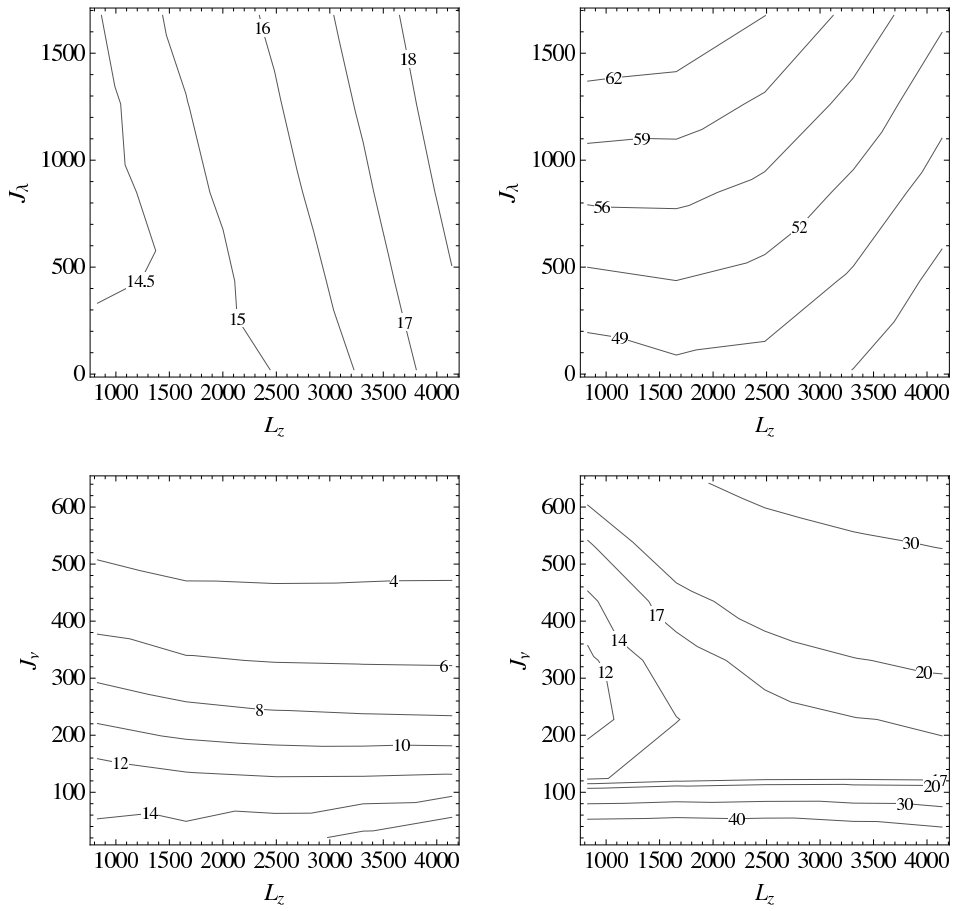

Figure 5.26: Details of stream geometry for the Stäckel potential SP1. Left panels: contours for the misalignment angle ϑ, in degrees, between the principal eigenvector of $\mathbf{D}$ and $\mathbf{\Omega}_0$, shown as a function of $\mathbf{J}$. Right panels: contours for the eigenvalue ratio λ_1/λ_2. The top panels show the plane in action-space with $J_\nu = 20.7\,\mathrm{kpc\,km\,s^{-1}}$, while the bottom panels show the plane in action-space with $L_z = 414\,\mathrm{kpc\,km\,s^{-1}}$. The range of actions covers a variety of interesting orbits: details of orbits at the extremes of the range are given in Table 5.6

while the right panels of Fig. 5.25 show contours of constant density and constant potential for this model.

5.7.4 Stream misalignment in Stäckel potentials

We now consider the geometry of streams formed in the Stäckel potentials SP1 and SP2. Although the form of the Stäckel potential allows the Hamilton-Jacobi equation to separate, and thus allows the actions $\mathbf{J}$ to be defined in terms of an integral over a single coordinate, expressions for $\mathbf{J}$ do not exist in closed form. Instead, the integrals in the expressions for $\mathbf{J}$ have to be evaluated numerically. Similarly, expressions for both the frequencies $\mathbf{\Omega}$ and their derivatives $\nabla_{\mathbf{J}}\mathbf{\Omega}$ can be written down, but not in closed form, and the integrals that they contain must too be evaluated numerically.

The details of the expressions for $\mathbf{J}$, $\mathbf{\Omega}$ and $\nabla_{\mathbf{J}}\mathbf{\Omega}$, and how to evaluate them, appear in Appendix A.

Table 5.6: The coordinate extrema of selected orbits from Fig. 5.26, illustrating the variety of orbits covered by that figure. The actions are expressed in $\mathrm{kpc\,km\,s^{-1}}$, while the apses are in kpc.

J_λ	L_z	J_ν	R_p	R_a	$\lvert z\rvert_{\max}$
20	828	20	3.5	5	0.74
20	4140	20	20	26	2.5
1680	4140	20	14	70	7
1680	828	20	1.75	27	3
414	828	20	2	10	1.25
414	4140	20	16	38	3.5
414	4140	640	22	56	34
414	828	640	3	20	17

Here, we simply note that having evaluated these quantities for a particular orbit, the eigenvectors $\hat{\mathbf{e}}_n$ and the eigenvalues λ_n are computed directly from the matrix $\mathbf{D}(\mathbf{J}_0) = \nabla_{\mathbf{J}}\Omega|_{\mathbf{J}_0}$ by standard methods.

We now examine the geometry of streams in the SP1 potential. The left panels of Fig. 5.26 show contour plots of the misalignment ϑ in angle-space between the principal direction of $\mathbf{D}$ and the frequency vector $\mathbf{\Omega}_0$, where ϑ is calculated from equation (5.31), as was the case for systems of two actions. The right panels of the same figure show contours of the ratio λ_1/λ_2. The range of actions shown in these plots covers a variety of interesting orbits; the apses of the orbits at the extremes of the range are described in Table 5.6

As with the equivalent plots for the isochrone potential (Fig. 5.5), we see that the principal direction of $\mathbf{D}$ is never perfectly aligned with $\mathbf{\Omega}_0$. We see that in this very flattened potential, those streams with low J_ν have the greatest degree of misalignment, at about $\sim 15^\circ$. These orbits spend much of their time near the disk, and never get very far from it. The misalignment diminishes with increasing J_ν, falling to $\sim 4^\circ$ for orbits with apses in z of some tens of kpc. Hence, in this very flattened potential, there is much more prospect for dramatic misalignment than with the isochrone potential, which Fig. 5.5 shows to cause only comparatively smaller misalignments. Fig. 5.26 also shows that the eigenvalue ratio $\lambda_1/\lambda_2 > 10$ everywhere for the SP1 potential; thus, we conclude that highly elongated streams will form on all orbits which permit a cluster to be disrupted.

Fig. 5.27 shows the equivalent plot to Fig. 5.26, but for the SP2 potential. Like with the SP1 potential, the principal eigenvector is everywhere misaligned with $\mathbf{\Omega}_0$. The misalignment is maximized for plunging orbits, and minimized for highly inclined orbits. However, the magnitude of the misalignment is everywhere much smaller than is observed in the SP1 potential. For orbits that remain close to the plane, the magnitude of the misalignment is comparable to or larger than that seen in the isochrone potential (Fig. 5.5). For highly inclined orbits, the misalignment is slightly less than is seen in the isochrone potential, for orbits with similar apses.

Fig. 5.27 also shows that, like in the SP1 potential, the ratio of the eigenvalues of $\mathbf{D}$ is everywhere

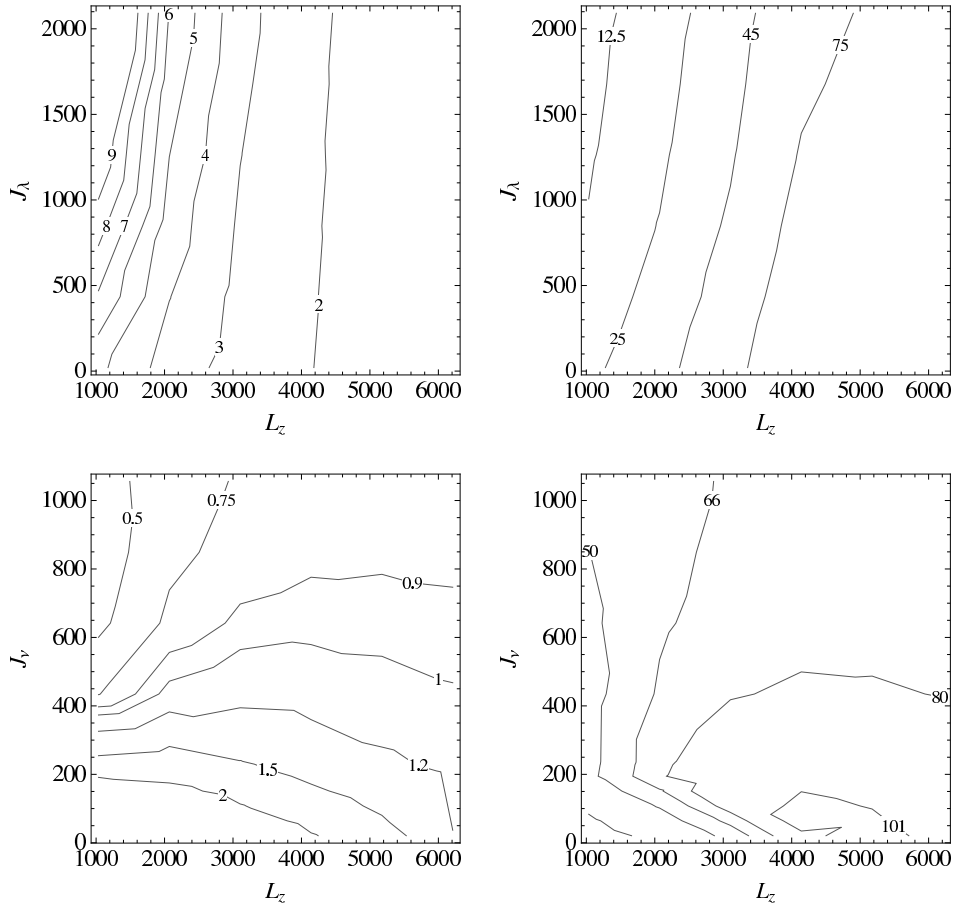

Figure 5.27: Similar to Fig. 5.26, but for the Stäckel potential SP2. The top panels show the plane in action-space with $J_\nu = 20.7\,\mathrm{kpc\,km\,s^{-1}}$, while the bottom panels show the plane in action-space with $L_z = 414\,\mathrm{kpc\,km\,s^{-1}}$. The range of actions covers a variety of interesting orbits: the behaviour of orbits at the extremes of the range is shown in Table 5.7

Table 5.7: The coordinate extrema of selected orbits from Fig. 5.27, illustrating the variety of orbits covered by that figure. The actions are expressed in $\mathrm{kpc\,km\,s^{-1}}$, while the apses are in kpc.

J_λ	L_z	J_ν	R_p	R_a	$\lvert z\rvert_{\max}$
20	1030	20	3.5	5	1.5
2100	1030	20	2	30	7
20	6200	20	27	31	3
2100	6200	20	15	65	10
410	1030	20	2.5	10	5
410	6200	20	21	40	7
410	1030	1050	3	22	20
410	6200	1050	22	50	32

Table 5.8: Actions and apses for selected orbits in the Stäckel potentials used in this chapter. Since all examples are in axisymmetric potentials, $J_\phi = L_z$ generally. The trajectories of these orbits are illustrated in Fig. 5.28.

	$J_\lambda/\,\mathrm{kpc\,km\,s^{-1}}$	$L_z/\,\mathrm{kpc\,km\,s^{-1}}$	$J_\nu/\,\mathrm{kpc\,km\,s^{-1}}$	R_p	R_a	$\lvert z\rvert_{\max}$
SO1	252.3	2618.	20.5	8	18	2
GD1	143.4	2871.	357.3	10	19.5	12
OS1	1502.	2391.	533.4	6	37	26

large. Hence, disrupted clusters should always form elongated streams in this potential.

In conclusion, we have found that in flattened Stäckel potentials with asymptotic logarithmic behaviour, streams will form from disrupted clusters on all realistic orbits, and that such streams will be generally misaligned with the orbits of the stars that compose them. If we take such potentials to be representative of the potential of our own Galaxy, we must conclude that, generally, streams observed in and around the Milky Way galaxy will not be perfectly aligned with orbits.

The precise behaviour of any given stream depends on both the potential and the action-space distribution of its stars. To proceed further we must again consider specific examples, by means of N-body simulation.

5.7.5 A stream in the Stäckel potential SP1

The top panels of Fig. 5.28 show the real-space trajectory of the orbit SO1 (Table 5.8) in the Stäckel potential SP1 (Table 5.5). This orbit has apses of approximately $R = (8, 18)\,\mathrm{kpc}$ in the galactic plane, and $z = (-2, 2)$ above and below the plane. It is thus fairly representative of an eccentric orbit that might be occupied by a globular cluster embedded in a galactic disk.

The cluster model C5 (Table 5.4) describes a King model specified for the orbit SO1 according to the schema of §5.6.1. The model has the same mass and profile parameter as does C1, and is very similar in all other attributes, because the orbit SO1 is not entirely dissimilar to the orbit I4 for which C1 was specified. A 10^4 particle realization of the C5 was made by random sampling of the King model distribution function. This cluster was placed close to apocentre on the orbit SO1 and evolved forward

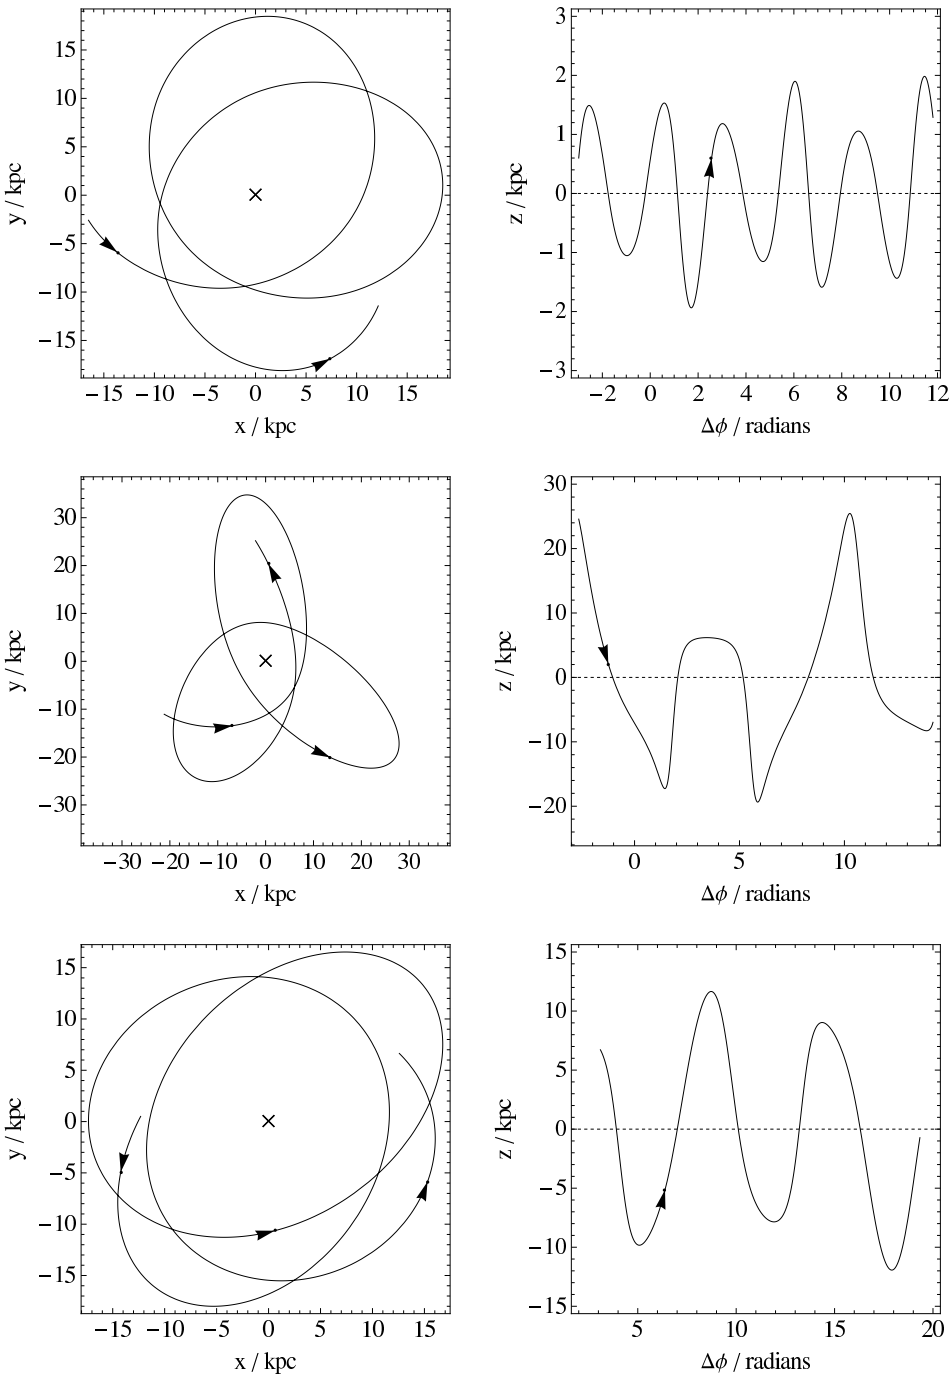

Figure 5.28: The real-space trajectories of the orbits SO1 (top panels), GD1 (middle panels) and OS1 (bottom panels), as described in Table 5.8. The trajectories shown were evaluated in the Stäckel potentials SP1 (for SO1) and SP2 (for GD1 and OS1), which are described in Table 5.5.

in time by the FVFPS tree code, with time step $\mathrm{d}t = t_{\mathrm{dyn}}/100$ and softening parameter ϵ as specified in Table 5.4. The total period of the simulation was 2.15 Gyr, or 7 complete radial oscillations.

5.7.5.1 Action-space distribution

Fig. 5.29 shows the evolution of the action-space distribution of this simulated cluster with time. Each row of panels shows the distribution at a different point in time. The left panel in each row shows the orthographic projection of the actions onto the (L_z, J_λ) plane, while the right panel of each row shows a similar projection onto the (L_z, J_ν) plane. In all panels, the appropriate projection of the mapped frequency vector, $\mathbf{D}^{-1}\boldsymbol{\Omega}_0$, is shown as an arrowed black line.

The top row shows the actions when the cluster is near to its first pericentre passage. In the left panel, the distribution is somewhat flattened, and oriented with positive gradient in $\Delta J_\lambda/\Delta L$. This behaviour is analogous with that seen in the top-right panel of Fig. 5.19: the motion of the cluster is predominantly in (λ, ϕ), thus (J_λ, L_z) are good proxies for the radial action J_r and angular momentum L respectively. J_λ and L_z are therefore highly correlated for a cluster near apsis on this orbit, in analogy with the mechanism described by equation (5.76) in §5.6.2.

Conversely, the distribution in the right panel, while being narrow, is oriented almost exactly along $\hat{L}_z$. We can understand the shape as follows. For this orbit, which is confined to be close to the plane, $J_\nu \sim J_z/2$, where the factor of 2 appears because J_ν is defined on a path restricted to only one side of the plane. J_z can be estimated by close analogy with equation (5.50). Hence, the spread in J_ν for a cluster of velocity dispersion σ is approximately

$$\Delta J_\nu \simeq \frac{1}{2}\Delta J_z \sim \frac{1}{2\pi}\delta p_z \Delta z \simeq \frac{1}{2\pi}\sigma \Delta z. \tag{5.86}$$

By analogy with equation (5.52) we find

$$\frac{\Delta J_\nu}{\Delta L_z} \sim \frac{\Delta z}{2\pi R_p}, \tag{5.87}$$

where R_p is the galactocentric pericentre radius in cylindrical coordinates. Evaluating this expression for the orbit SO1 gives $\Delta J_\nu/\Delta L_z \sim 0.04$, which we see from the top-right panel of Fig. 5.29 is close to exact.

The flat orientation of the top-right panel we explain by pointing out that, as the top panels of Fig. 5.28 show, the motion in ν in this example is almost decoupled from the radial motion. This means that the ν coordinate need not be at apsis when the cluster is at pericentre, and thus the arguments of §5.6.2, which force a correlation between $J_\lambda \sim J_r$ and $L_z \sim L$ near pericentre, do not apply. For an orbit

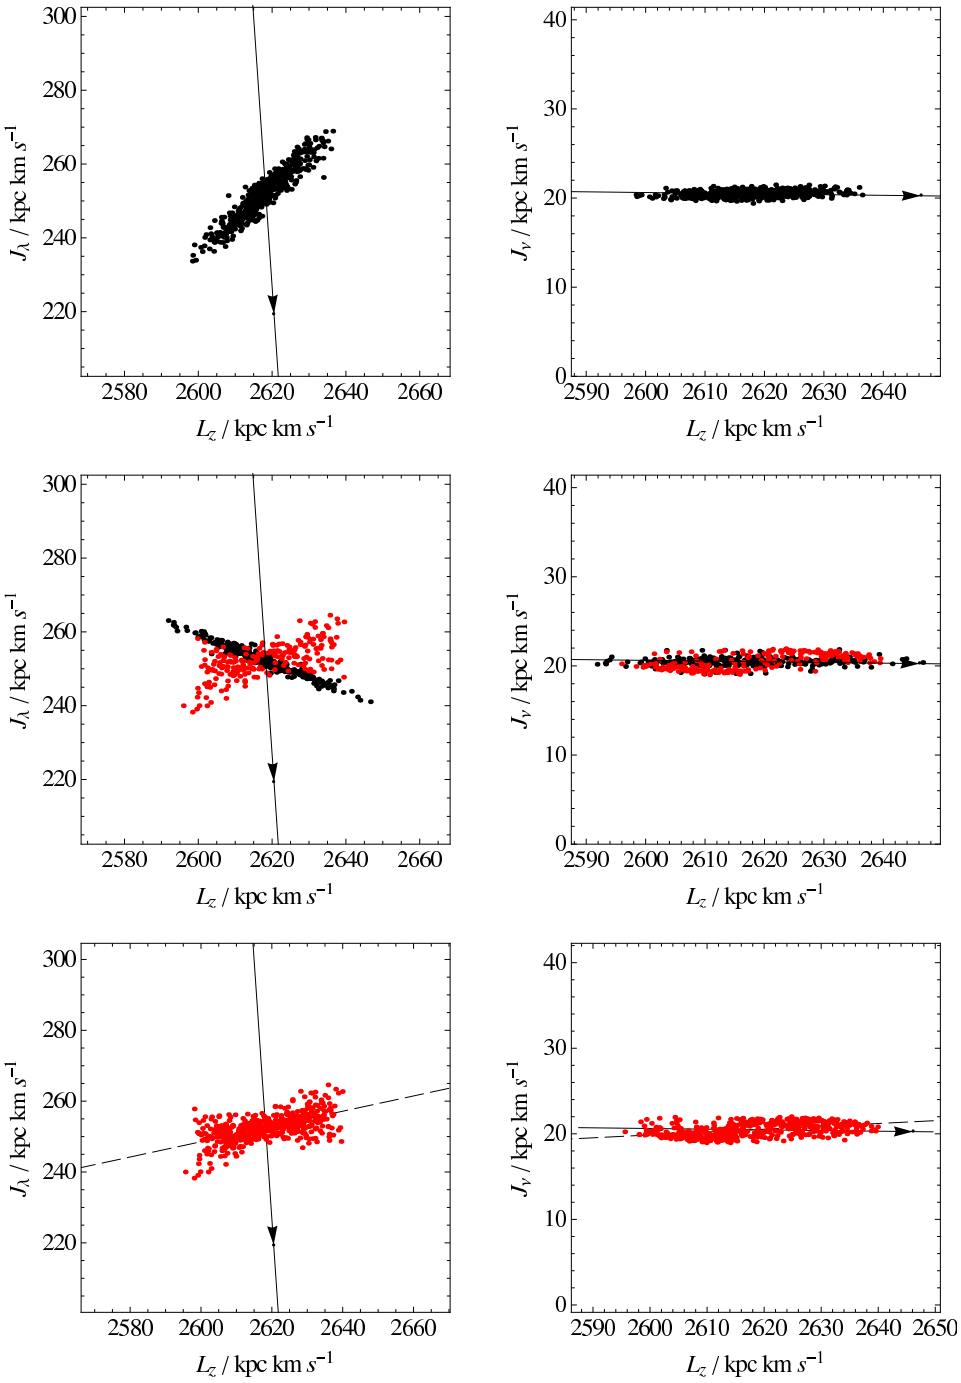

Figure 5.29: Action-space distribution for the simulated cluster C5, on the orbit SO1 in the Stäckel potential SP1, at various points in time. The left panels show the orthographic projection of distribution onto the (L_z, J_λ) plane, while the right panels show a similar projection onto the (L_z, J_ν) plane. Distributions are shown at the following times: top panels, the first pericentre passage; middle panels, the subsequent apocentre passage; bottom panels, the 7th apocentre passage. Black particles are bound to the cluster, while red particles orbit free in the host potential. The dashed lines in the bottom panels are lines that have been least-squares fitted to the particles.

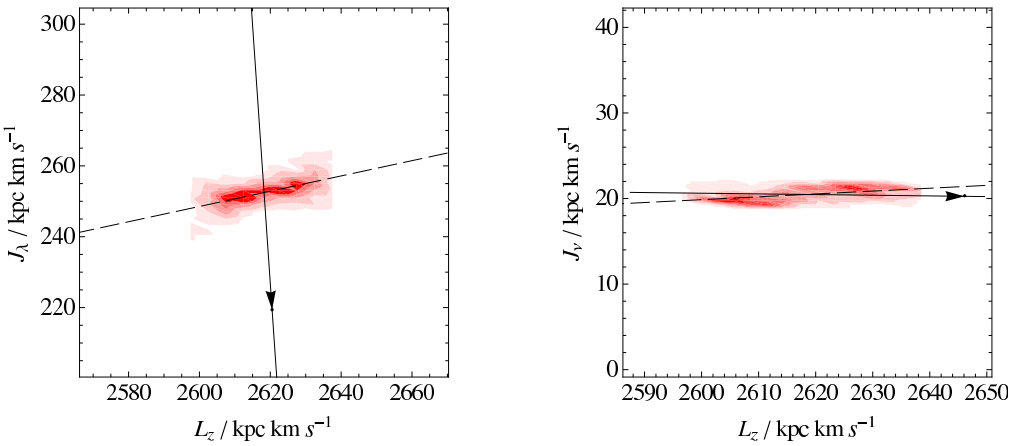

Figure 5.30: Plots showing the (column) density of particles in action-space, corresponding to (left) the bottom-left and (right) the bottom-right panels of Fig. 5.29. The density was estimated by placing particles into bins of width $\sim 2\,\mathrm{kpc\,km\,s^{-1}}$. Darker shading represents regions of higher particle density, with the edges of the shaded regions representing contours of constant density. As was true of Fig. 5.22, no significant sub-structure in the density field is visible that is not attributable to sampling noise.

in which J_ν is more strongly coupled to the radial motion, we would expect to see the characteristic tilting of the (L_z, J_ν) distribution near pericentre and apocentre, as a correlation between J_ν and L is forced.

The middle panels of Fig. 5.29 show the action-space distribution at the subsequent apocentre passage. Bound and unbound particles are shown in black and red, respectively. The action-space structure of the unbound particles in the left panel bears striking similarity to that shown in Fig. 5.19, as might be expected when (J_λ, L_z) make good proxies for (J_r, L). The same physical principle for the disruption of the cluster applies here as it does in the spherical case; that is, the particles will escape the cluster through the Lagrange points L_1 and L_2 by first travelling radially, so the action-space distribution will be compressed in this direction. Thus, the range of ΔL_z and ΔJ_ν for the unbound stars is about the same as in the top panels, but the range of ΔJ_λ is markedly less.

The bottom panels of Fig. 5.29 show action-space distribution at the 7th apocentre passage. The structure is essentially the same as that of the middle panels, except that all particles are now unbound. Also plotted in each of the bottom panels is a dashed line, which has been least-squares fitted to the unbound particles. We note that the image of the frequency vector and the dashed line are highly misaligned in the bottom-left plot; hence, we expect the stream to be significantly misaligned with $\mathbf{\Omega}_0$ in angle-space.

In order to better examine the variation in particle density within the bottom panels of Fig. 5.29, Fig. 5.30 shows equivalent plots of particle density. As was seen in Fig. 5.22, the plots reveal the particles to be primarily concentrated in the lobes of the distribution, but they fail to reveal any significant

additional structure in the particle-density field.

In conclusion, we find that in very flattened potentials, disrupted clusters form an action-space distribution that is wholly analogous with that found for disrupted clusters in spherical potentials. In the following section, we test the ability of a simple straight-line model of the action-space distribution to predict the track of the stream.

5.7.5.2 The effects of disk shocks

Unlike with a spherical potential, an axisymmetric potential allows for tidal forces other than those felt during pericentre passage to act upon an orbiting cluster. In particular, the passage of a cluster through a massive galactic disk will subject a cluster to a tidal force that is of comparable magnitude to that felt when close to pericentre.[9]

The tidal stress imposed on a cluster at pericentre has a tensile component, which acts to strip stars from the cluster. Conversely, the tidal stress imposed by a disk passage is entirely compressive in nature. Hence, stars are not actively stripped from a cluster during a disk passage. Instead, the action of such 'disk shocks' is to heat the cluster, perhaps repopulating the outer edges of the cluster, the stars from which were stripped during a previous pericentre encounter (Spitzer, 1987, §5.2a).

The net effect of disk shocks on the stripping process is a faster and more complete disruption of the cluster than would take place for an unshocked cluster exposed to equivalent pericentric tidal stress. Since the vast majority of stars continue to be stripped at pericentre even when the effect of disk shocks is significant, the gross action-space distribution resulting from the stripping of a shocked cluster will remain as previously described. However, since the disk shocks act to increase the velocity dispersion of the cluster between stripping events, it is likely that the wings of the resulting action-space distribution will be populated with more stars than would otherwise have been the case.

5.7.5.3 Predicting the stream track

Fig. 5.31 shows the angle-space configuration for the simulated cluster near its 7th apocentre passage. The grey particles are for angles that have been computed directly from the output of the N-body simulation. The arrowed black line is $\mathbf{\Omega}_0$, while the blue line is the map of the dashed-line from Fig. 5.29. In both

[9] To see this, we note that the tidal force in the direction $\hat{\mathbf{x}}$ at any given point,

$$\frac{\mathrm{d}\mathbf{F}}{\mathrm{d}x} = (\hat{\mathbf{x}} \cdot \nabla)\,\mathbf{F} = -(\hat{\mathbf{x}} \cdot \nabla)\,\nabla\phi \sim 4\pi G\rho, \tag{5.88}$$

is proportional to ρ, which we understand to be the mean density of tide-inducing matter interior to the point under consideration. Consider then the example of a Milky Way cluster undergoing both pericentre passage, and encountering disk shocks, at a galactocentric radius of $R_0 \sim 8\,\mathrm{kpc}$. The disk density in the solar neighbourhood is $\rho_d \sim 0.1 M_\odot\,\mathrm{pc}^{-3}$ (Binney & Tremaine, 2008, Table 1.1). Meanwhile, the mean density interior to the Sun required to give a circular speed of $v_c \sim 240\,\mathrm{km\,s}^{-1}$ at R_0 is $\rho_g \sim 0.06 M_\odot\,\mathrm{pc}^{-3}$. Hence, the magnitude of the tidal effects felt during disk passage is indeed comparable to those felt during pericentre passage.

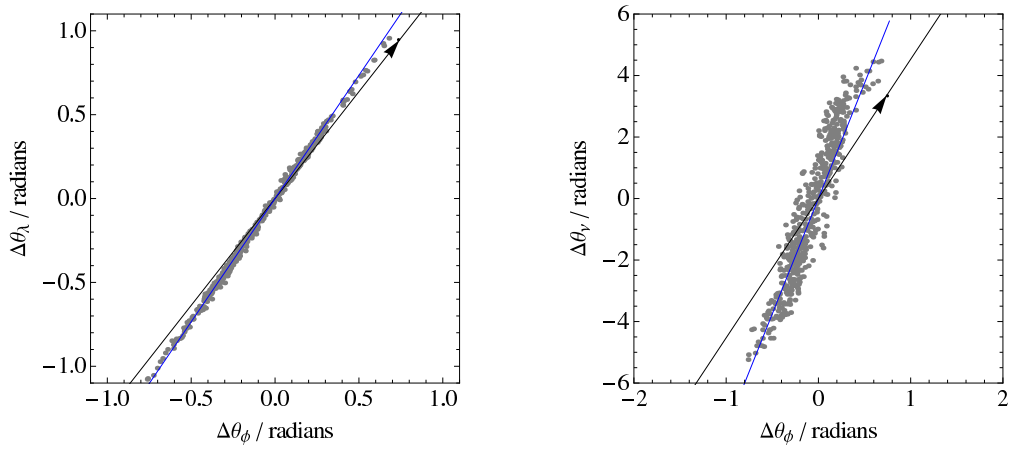

Figure 5.31: Angles for the N-body cluster shown in the bottom panels of Fig. 5.29. The blue line shows the predicted stream, resulting from the mapping of the dashed line in bottom panels of Fig. 5.29. The blue line is clearly a much better representation of the stream than is $\mathbf{\Omega}_0$, represented by an arrowed black line.

projections, the blue line is clearly a much superior match to the data than is the orbit.

Fig. 5.32 shows equivalent plot to Fig. 5.31, but in real-space. The misalignment between the stream and the progenitor orbit in angle-space is seen to map into a large misalignment in real-space. Attempting to constrain halo parameters by fitting orbits to the stream shown in Fig. 5.32 would not produce sensible results. Conversely, the map of the dashed-line model for the action-space distribution, shown in Fig. 5.29, clearly provides an excellent proxy for the track of the stream.

We therefore conclude that, in flattened systems, we are able to accurately predict the track a of stream using a simple, but well informed, action-space model, while the corresponding progenitor orbit makes a substantially less accurate proxy for the track of a stream.

5.7.6 Realistic examples in the Stäckel potential SP2

In this section, we draw on the work of the previous sections to construct N-body simulations for two actual observed Milky Way streams, in order to examine to what extent these streams can act as proxies for progenitor orbits.

In the following sections, we take as our model of the Milky Way the Stäckel potential SP2, described in Table 5.5. This model does reproduce an approximately correct in-plane rotation curve outside of the solar circle, precisely where our example streams reside. However, this model is arguably insufficient in representing the flattening of the potential in proximity to the Milky Way's massive disk. §5.7.4 showed that stream misalignment is likely to worsen in the presence of a flatter potential. Thus, if our results below are in error, it is likely that streams make poorer proxies for orbits, not better ones.

Unfortunately, there is insufficient flexibility in Stäckel models to allow for an accurate representation

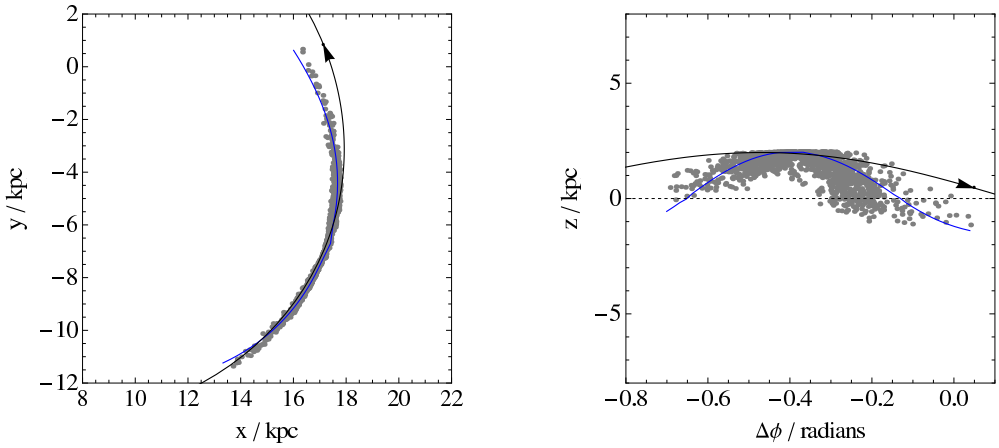

Figure 5.32: Real-space representation of an N-body simulation of cluster model C6 evolved on the orbit SO1. The potential is the flattened Stäckel potential of Table 5.1. The points are shown near the 7th apocentre passage following release. The black line shows the trajectory of an individual orbit; the blue line shows the predicted stream path.

of both halo and disk across a reasonable range of radii. Correcting for any error due to insufficient flattening, or even properly assessing the magnitude of its effects, is therefore beyond the scope of this chapter. It is with this caveat in mind that we now proceed.

5.7.6.1 Tidal stream GD-1

The tidal stream GD-1 (Grillmair & Dionatos, 2006) is associated with one of the most complete sets of full phase-space observations for any stream in the Milky Way, and has been utilized in several recent attempts (Willett et al., 2009; Koposov et al., 2010) to constrain the Galactic potential by fitting an orbit to it. Furthermore, it is the example used by this author (Chapter 4; Eyre, 2010) to demonstrate the applicability of Galactic parallax, the analysis of which would be subject to systematic error if the track of the stream were not tangent with the orbits of the stars. It is therefore of primary importance that we understand the behaviour of this particular stream.

The orbit GD1, described in Table 5.8 and illustrated in the middle panels of Fig. 5.28, was chosen to approximate the observed features of the GD-1 stream. In particular, the orbit was required to follow the conclusions of Chapter 4: i.e. it should be at pericentre, about 7 kpc distant from our fiducial Solar location; it should have an orbital plane inclined to the Galactic plane by 37°, and be on a retrograde orbit; and it should reproduce approximately the track on the sky as reported in Koposov et al. (2010).

Having selected an orbit in the SP2 potential that roughly meets these requirements, the cluster model C6 was specified according to the schema of §5.6.1. In the absence of any evidence as to the true nature of the GD-1 progenitor, the model was chosen to have the same mass and profile parameters as does C1. The stripping radius $r_s = 15$ kpc was chosen to be slightly larger than the pericentre radius of

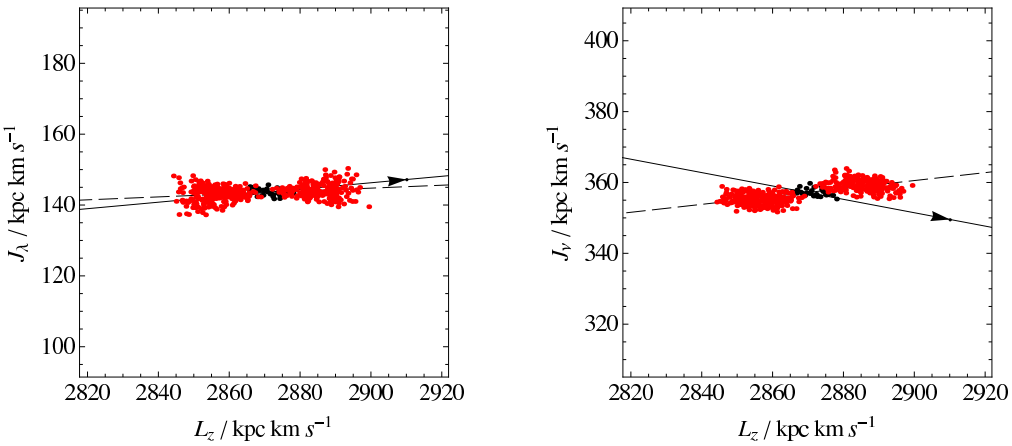

Figure 5.33: Actions for the cluster C6 on orbit GD1 in the Stäckel potential SP2. The result is intended to simulate a possible configuration of the GD-1 Milky Way stream. The cluster is shown near pericentre after $14\frac{1}{2}$ radial orbits. The arrowed black line is the image of the frequency vector, $\mathbf{D}^{-1}\mathbf{\Omega}_0$; the dashed line is a straight-line fit to the data.

the GD1 orbit.

A 10^4 particle realization of the C6 model was made by random sampling of the King model distribution function. This cluster was then placed close to apocentre on the orbit GD1, at a point 4.57 Gyr prior to the present pericentre location, equivalent to $14\frac{1}{2}$ radial orbits. The cluster was then evolved forward in time by the FVFPS tree code, with time step $\mathrm{d}t = t_{\mathrm{dyn}}/100$ and softening parameter ϵ as specified in Table 5.4.

Fig. 5.33 shows the action-space configuration of the simulated GD-1 cluster at the end of the simulation period, with the cluster centroid at pericentre, and the stream near its present location. Bound and unbound particles are plotted in black and red, respectively; we note that the cluster has been almost completely disrupted at this time. The image of the frequency vector, $\mathbf{D}^{-1}\mathbf{\Omega}_0$, is shown in this figure as an arrowed black line, while the dashed line is a least-squares fit to the unbound particles. In addition, Fig. 5.34 shows particle-density plots corresponding to the scatter plots of Fig. 5.33.

The action-space distribution does not look dissimilar to those in either Fig. 5.23 or Fig. 5.29. In this eccentric, highly inclined orbit, the ellipsoidal coordinates (λ,ν) effectively parametrize the spherical polar coordinates (r, θ). The action J_λ can therefore be compared to the radial action J_r, with J_ν becoming some function of the total angular momentum L and L_z. In this example, J_ν is comparatively small, so $L \sim L_z$ and we can understand the left panel of Fig. 5.33 by direct analogy with the results of §5.6.2 and §5.7.5.1.

We will not attempt to disentangle the angular momentum L, in order to make sense of the right panel. Instead, we merely note that in comparison to the previous examples, the action-space image of the frequency vector $\mathbf{D}^{-1}\mathbf{\Omega}_0$ is aligned much more closely with the long axis of the action-space

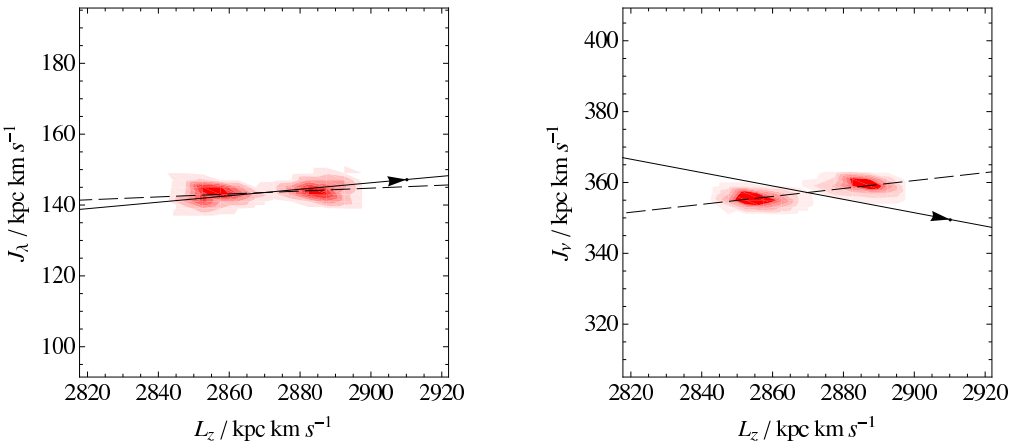

Figure 5.34: Plots showing the (column) density of particles in action-space, corresponding to the scatter plots of Fig. 5.33. The density was estimated by placing particles into bins of width $\sim 2\,\mathrm{kpc\,km\,s^{-1}}$. Darker shading represents regions of higher particle density, with the edges of the shaded regions representing contours of constant density.

distribution. We therefore expect the misalignment between the angle-space stream and the frequency vector to be less here than in previous examples.

We can actually understand this alignment as a property of spherical, logarithmic potentials, which is approximately what this distant orbit feels in the Stäckel potential SP2. In such a potential, where

$$\Phi(r) = v_c^2 \log r/r_0, \tag{5.89}$$

the circular frequency Ω_ϕ is given by

$$\Omega_\phi = \frac{v_c^2}{L}. \tag{5.90}$$

To first order, the radial frequency Ω_r is given by (Binney & Tremaine, 2008, §3.2.3)

$$\Omega_r^2 = \left.\frac{\partial^2 \Phi_{\mathrm{eff}}}{\partial r^2}\right|_{r_g} = \left.\frac{\partial^2 \Phi}{\partial r^2}\right|_{r_g} + \frac{3L^2}{r_g^4} = -\frac{v_c^2}{r_g^2} + \frac{3v_c^4}{L^2} = \frac{2v_c^4}{L^2} = 2\Omega_L^2. \tag{5.91}$$

Although the above equation is only true to first order—and indeed, it is the higher-order terms that contribute to $\mathbf{D}$ and give rise to the stream geometry seen in Fig. 5.27—we see that

$$\mathbf{\Omega}_0 = (\Omega_r, \Omega_\phi) \sim \Omega_\phi(\sqrt{2}, 1) \simeq \frac{v_c^2}{L}(\sqrt{2}, 1). \tag{5.92}$$

The direction of $\mathbf{\Omega}_0$ is approximately constant, and the magnitude is a function of L only. Hence, the image of $\hat{\mathbf{\Omega}}_0$ in action-space points almost exactly along $\hat{L}$.

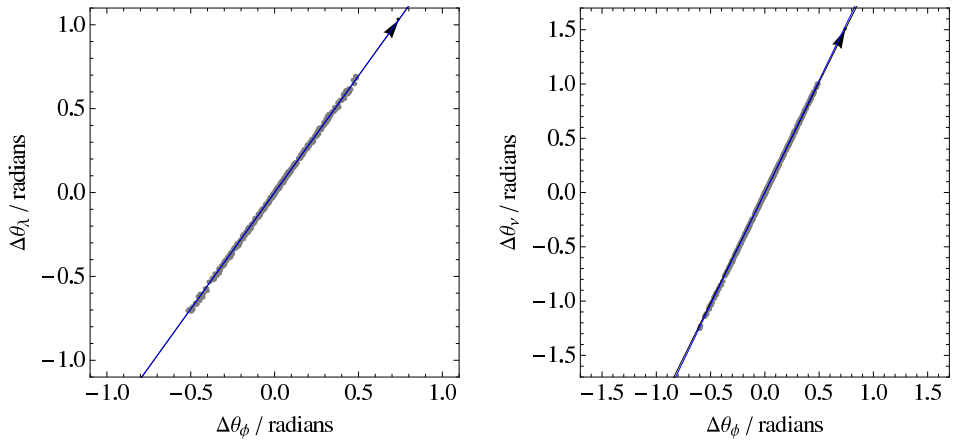

Figure 5.35: Angle-space configuration of the simulated GD-1 stream, shown in Fig. 5.33. The grey particles are computed directly from the results of the N-body simulation, while the blue line is the angle-space map of the dashed line from Fig. 5.33. The arrowed black like shows the frequency vector, $\mathbf{\Omega}_0$.

We have seen from the arguments in §5.6.2 that the action-space distributions of clusters formed in spherical potentials are likely to be flattened and oriented towards $\hat{L}$. Thus, in spherical logarithmic potentials, we expect the long axis of such distributions to be aligned closely with the action-space image of the frequency vector. Correspondingly, we expect streams in such potentials to be more closely aligned with $\mathbf{\Omega}_0$ in angle-space than was seen in the isochrone potential or in the highly-flattened SP1 potential.

Returning to our example, Fig. 5.35 shows the angle-space configuration of the simulated GD-1 stream at the end of the simulation period. The grey particles are for angles that were computed directly from the output of the N-body simulation. The arrowed black line is $\mathbf{\Omega}_0$, while the blue line is the angle-space map of the dashed line from Fig. 5.33. As expected given the close alignment of the action-space distribution with the image of $\mathbf{\Omega}_0$, the stream, the frequency vector and the predicted track are all seen to agree almost perfectly.

Fig. 5.36 shows the real-space configuration of the simulated GD-1 stream and the end of the simulation period. The plots have been rendered in Galactic coordinates (l, b) and heliocentric distance $r_{\rm sun}$ to aid with comparison to observations. The GD-1 stream is seen to be an excellent proxy for the progenitor orbit. This is expected, given the perfect alignment of the stream with $\mathbf{\Omega}_0$ in Fig. 5.35. A small discrepancy of $\sim 0.3\,$kpc does exist between the heliocentric distance of the stream and the orbit at the extreme end of the leading tail, but the match is otherwise perfect. In the case of the on-sky projection, the stream perfectly delineates the orbit for all longitudes. In both panels, the blue line, which is the track predicted from the dashed line in Fig. 5.33, is a perfect match to the stream everywhere.

In conclusion, we find that the simulated GD-1 stream delineates its orbit perfectly. This is fortunate for our analysis of the Galactic parallax of the GD-1 stream in Chapter 4, which does not now need

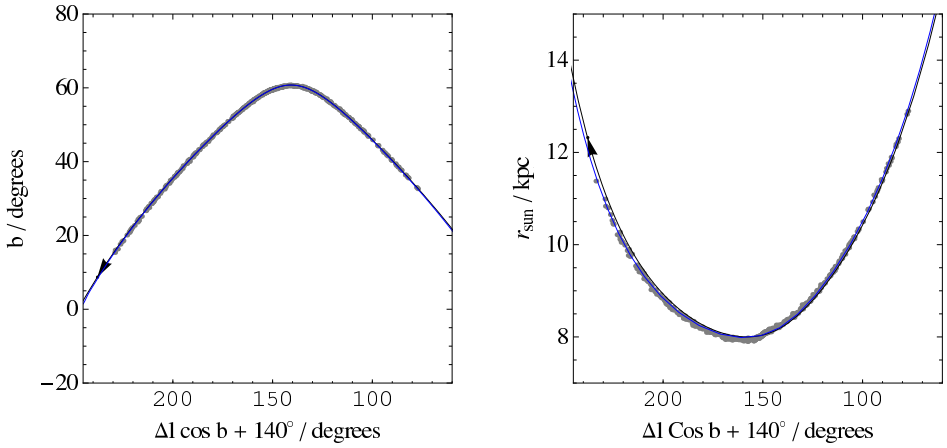

Figure 5.36: Real-space configuration of the simulated GD-1 stream, shown in Fig. 5.33. Plots are shown in Galactic latitude/longitude and heliocentric radius, to aid comparison with observations. Left panel: the on-sky projection of the stream, which is seen to agree perfectly with the progenitor orbit. Right panel: heliocentric distance, versus Galactic longitude. In the right panel there exists a tiny anomaly between the heliocentric distance of the stream and the trajectory of the orbit on the far left of the plot, but the agreement between stream and orbit is otherwise excellent.

to be revisited. We have understood the result in terms of a peculiar property of spherical logarithmic potentials, which causes the $\hat{L}$ direction in action-space to be mapped close to $\mathbf{\Omega}_0$ in angle-space. Since action-space distributions are naturally oriented close to $\hat{L}$, the result is an angle-space stream closely aligned with $\mathbf{\Omega}_0$.

5.7.6.2 Orphan stream

Our investigation into the utilization of tidal streams began with the Orphan stream (Belokurov et al., 2007), and the observation that standard techniques (e.g. Fellhauer et al., 2007b) had difficulty in finding an orbit that described it exactly. It is thus fitting that we end with the Orphan stream as our final case-study in this thesis.

The orbit OS1, described in Table 5.8 and illustrated in the bottom panels of Fig. 5.28, is based on the simulated Orphan stream shown in Chapter 2, which in turn is based on the tentative velocity data presented in Belokurov et al. (2007). The orbit OS1 roughly reproduces the on-sky track, distances, and velocities from Belokurov et al. (2007) when integrated in the SP2 potential. Recently, a newer analysis of the Orphan stream by Newberg et al. (2010) utilizing SEGUE spectra (Yanny et al., 2009) has produced radial velocity data that are considerably more certain, and unfortunately inconsistent with, the Belokurov et al. (2007) radial velocity data points. Nonetheless, we still present our simulation of the Orphan stream, in order to dispel any notion based on the previous section, that tidal streams in approximately spherical, logarithmic potential must always perfectly delineate orbits.

The cluster model C7 was specified according to the schema of §5.6.1, for the orbit OS1 in the SP2

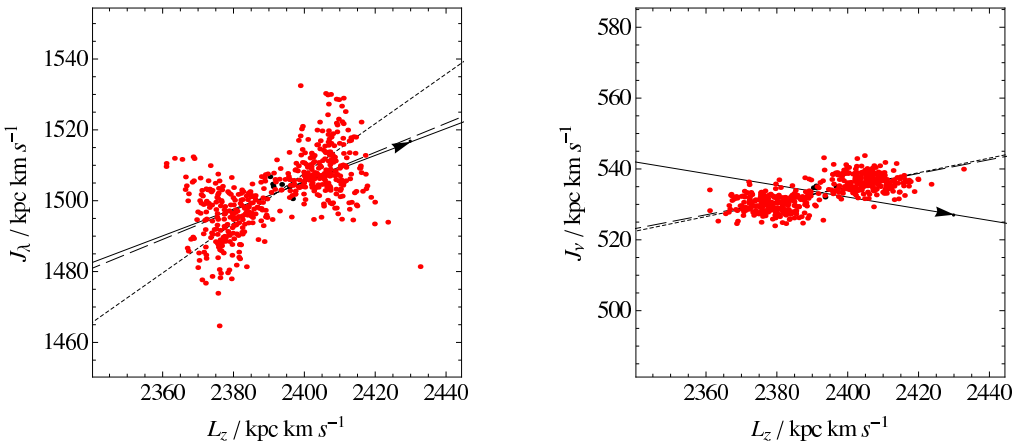

Figure 5.37: Actions for the N-body cluster C7 on orbit OS1 in the Stäckel potential SP2. The simulation is one possible configuration for the Orphan stream. The cluster is shown near apocentre after 14 complete radial orbits. The arrowed black line is the image of the frequency vector, $\mathbf{D}^{-1}\mathbf{\Omega}_0$. The dashed line is a least-squares fit to the data, and the dotted line is a weighted least-squares fit to the data, discussed in the text.

potential. Like with GD-1, evidence as to the true properties of the Orphan stream progenitor is lacking. Hence, the model was chosen to have the same mass and profile parameters as does C1; the stripping radius $r_s = 12\,\mathrm{kpc}$ was chosen to be slightly larger than the pericentre radius of the orbit.

A 10^4 particle realization of the C7 model was made. This cluster was then placed close to apocentre on the orbit OS1, at a point 6.52 Gyr prior to the present apocentre location, equivalent to 14 complete radial orbits. The cluster was then evolved forward in time by the FVFPS tree code, with time step $\mathrm{d}t = t_{\mathrm{dyn}}/100$ and softening parameter ϵ as specified in Table 5.4.

Fig. 5.37 shows scatter plots for the action-space distribution of the model Orphan stream at the end of the simulated period, while Fig. 5.38 shows the corresponding particle-density plots. The distribution is immediately comparable to that of Fig. 5.33, except that the structure in the left panel appears tilted when compared to that of GD-1. We understand this to result from a combination of two factors. First, the orbit OS1 is more eccentric than is the orbit GD1. From equation (5.76) we expect this to result in a distribution that is steeper when viewed in (L, J_r). Secondly, unlike in the two previous examples, a large component of L is in other than the z-direction. Hence, a distribution that is approximately aligned with $\hat{L}$ will not be aligned with $\hat{L}_z$. Again, we will not attempt to disentangle L in order to precisely understand the mechanics of the plots, but we note once more that the image of the frequency vector $\mathbf{D}^{-1}\mathbf{\Omega}_0$ and the long axis of the action-space distribution are closely aligned. We thus expect the angle-space stream to be well represented by $\mathbf{\Omega}_0$.

Also shown in Fig. 5.37 are two least-squares fitted lines: the dashed curve is an unweighted fit to the particles, while the dotted line is a weighted fit to the particles, described in detail below.

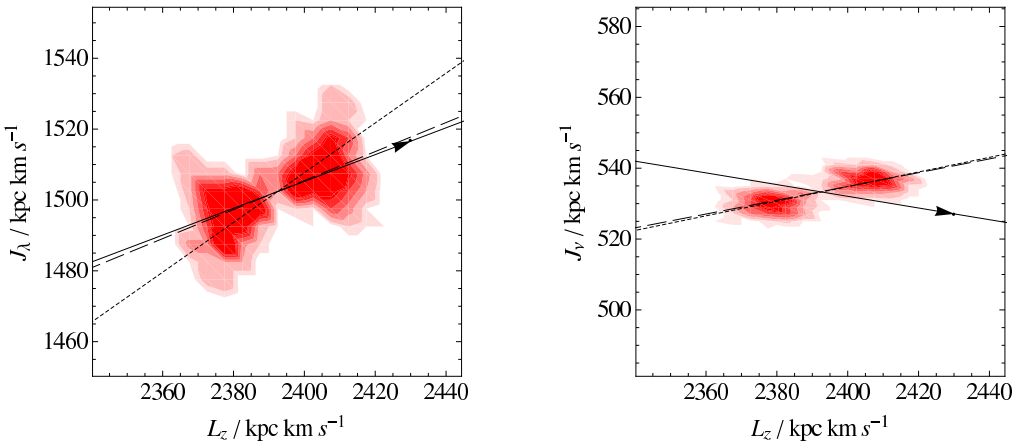

Figure 5.38: Plots showing the (column) density of particles in action-space, corresponding to the scatter plots of Fig. 5.37. The density was estimated by placing particles into bins of width $\sim 2\,\mathrm{kpc\,km\,s^{-1}}$. Darker shading represents regions of higher particle density, with the edges of the shaded regions representing contours of constant density.

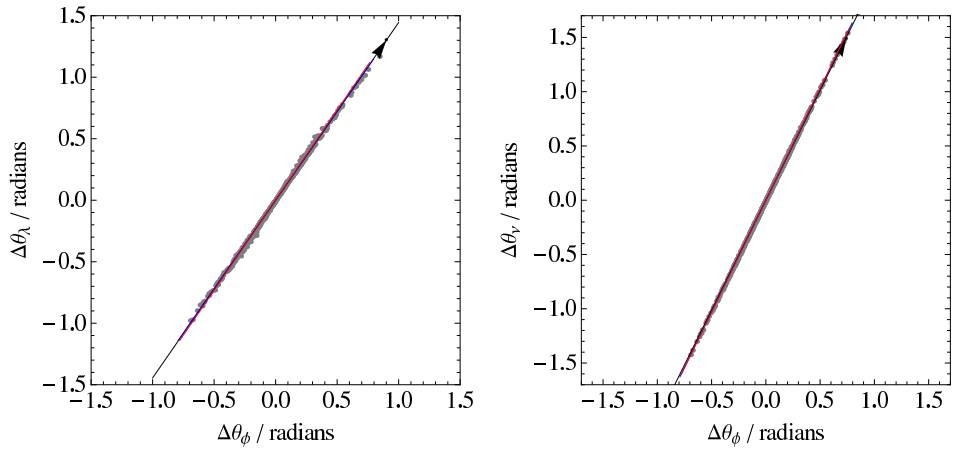

Figure 5.39: Angle-space configuration of the simulated Orphan stream, corresponding to the actions shown in Fig. 5.37. The angles of the grey particles have been computed directly from the N-body simulation. The arrowed black line is $\mathbf{\Omega}_0$, the blue line is the image of the dotted line from Fig. 5.37, and the red line is the image of the dashed line from Fig. 5.37.

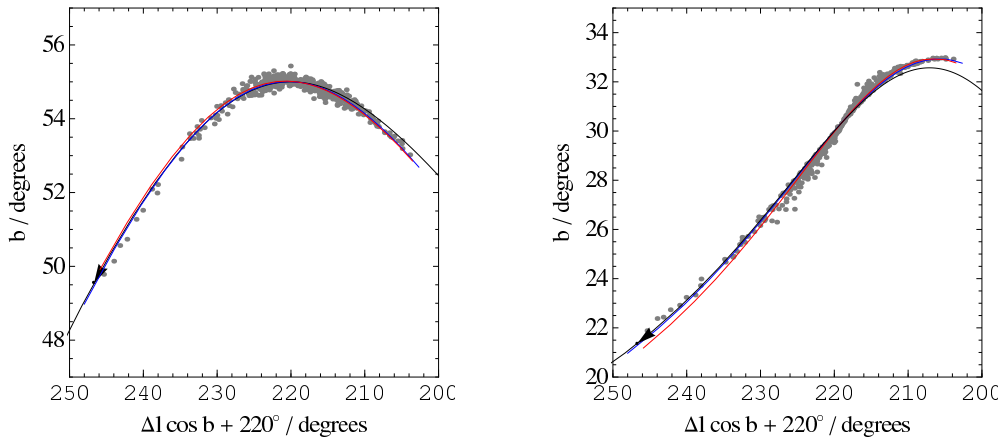

Figure 5.40: Real-space configuration of the simulated Orphan stream. The N-body particles are plotted in grey. The black line is the trajectory of the progenitor's orbit, while the blue and red lines are the real-space mappings of the blue and red lines from Fig. 5.39. There is a clear anomaly between the orbit and the stream on the right-hand side of both plots. On the left-hand side of the right panel, there is an anomaly between the red curve and the stream.

Fig. 5.39 shows the angle-space configuration of the simulated Orphan stream at the end of the simulation run. The grey particles are for angles that were computed directly from the output of the N-body simulation. In both panels $\mathbf{\Omega}_0$, shown as a black arrowed line, is almost indistinguishable from the red and blue lines, which are the angle-space images of the dashed and dotted lines, respectively, from Fig. 5.37. All three lines almost perfectly delineate the stream, although close inspection shows the red line to pass on either side of the particles at the extremes of the tail, while $\mathbf{\Omega}_0$ and the blue line pass directly through.

Fig. 5.40 shows the real-space configuration of the simulated stream, at the end of the simulation run. The grey particles are plotted directly from the output of the simulation, while the lines are the real-space equivalents of those shown in Fig. 5.39.

We immediately note that the stream and the orbit trajectory do not coincide on the right side of each plot. Thus, the stream is not well represented by an orbit in this region, despite the stream and $\mathbf{\Omega}_0$ coinciding everywhere in Fig. 5.39. The explanation is that, although there is no misalignment between the stream and the orbit in angle-space, the changes in trajectory induced by the variation in action down the stream have caused the stream and the orbit to become misaligned in real-space. The detail of this mechanism was described in §5.4.4. In mapping from angle-space to real-space, the blue and red curves include the correction for this effect, as specified in equation (5.53). Thus, the predicted tracks from the red and blue curves match the stream much better than does the orbit, on the right-hand side of each plot, where the magnitude of this effect is maximized.

Looking again at Fig. 5.40 we see that on the left side of the right panel, the red curve is a slightly

poorer representation of the stream than is the orbit or the blue curve. This occurs because the dashed-line model to the action-space distribution, shown in Fig. 5.37, is generating the wrong $\delta\mathbf{J}$ correction when utilized in equation (5.53). The wrong correction is generated because the stars at the ends of the tidal tails, rather than corresponding to the middling regions of the action-space distribution which are well modelled by the dashed line, actually correspond to the upper-right and lower-left extremities of the action-space distribution, which are well removed from the dashed-line model. Hence in this example, the dashed-line model does not provide an appropriate $\delta\mathbf{J}$ correction for the ends of the stream, and correspondingly the red line in Fig. 5.40 fails to properly predict the stream track at the end of the leading tail.

The remedy is a slightly more sophisticated model of the action-space distribution. The dotted line in Fig. 5.37 is a least-squares fit to the data, like the dashed line, but with each data point weighted by $\sigma_i = \mathbf{D} \cdot \Delta\mathbf{J}_i$. That is, we grant greater weight to those data points which map to the fastest diverging stars in the stream. The resulting model, though still simple, better predicts the variation in action down the stream, when utilized in equation (5.53), than does the dashed line. The results of mapping the dotted line into real-space are shown as the blue line in Fig. 5.40. Unlike the red line or the orbit, the blue line matches the stream track everywhere.

In conclusion, we find that the action-space distribution of a simulated Orphan stream exhibits the same characteristics as those observed in both the isochrone potential, and for other orbits in Stäckel potentials. We find that, although the principal direction of $\mathbf{D}$ is misaligned with $\mathbf{\Omega}_0$, the stream in angle-space is nonetheless aligned with $\mathbf{\Omega}_0$, because it is a property of spherical, logarithmic potentials that the $\hat{L}$ axis in action-space maps to $\mathbf{\Omega}_0$ in angle-space, and it is a natural feature of disrupted clusters that their action-space distribution is also aligned with $\hat{L}$. Conversely, we find that the stream is not well represented by the progenitor orbit at all points in real-space, because the variation in action of the stars as we move down the stream causes sufficient change to the trajectory of those stars, that the stream track and the progenitor orbit diverge.

We find that we can predict the path of the stream perfectly, but that the simple model of a line least-squares fitted to the action-space distribution is not sufficient to do so. An improved model, in which the data points are weighted by the quantity $\mathbf{D} \cdot \Delta\mathbf{J}$ for the least-squares fit, results in an action-space model that accurately predicts the stream track everywhere.

5.8 Conclusions

In this chapter, we have studied in detail the mechanics of the disruption of star clusters in various model potentials. We have been interested in learning of the conditions required for the formation of

tidal streams from such clusters, and in particular, to what extent such streams delineate the orbits of their stars. With regard to the latter, we have been motivated by the catalogue of techniques which attempt to utilize tidal streams to place constraints on the Galactic potential, many of which assume that streams perfectly delineate orbits.

We utilize action-angle variables extensively in our approach to the problem. It is found that these coordinates allow a convenient and natural description of the physical processes that occur in cluster disruption and stream formation. In the absence of more advanced techniques (e.g. McMillan & Binney, 2008), this approach has restricted the galaxy models under consideration to those for which action-angle variables are readily computed. Nonetheless, we do expect our findings to generalize to more sophisticated models of the Galactic potential.

In the broadest strokes, we have found that tidal streams will always form when a cluster in a realistic orbit around a realistic host galaxy is subject to tidal stripping. We show that these streams need not be well representative of an orbit in the host galaxy, and that in general they will not be, although clusters on particular orbits in particular potentials can well make an accurate representation. Particularly alarming is our finding that serious systematic errors can be made by attempting to constrain the Galactic potential using a stream that does not delineate an orbit. It is of some relief, however, that we find ourselves able to accurately predict the tracks of streams, even when they are not representative of an orbit, and so it may be possible to repair the potential-constraining procedures, by having them fit such stream tracks instead.

A complete summary of our findings in this chapter is presented below. The section that follows discusses some of the consequences of our findings, and presents some directions for future work. In Chapter 6 of this thesis we further review our findings in the context of contemporary Galactic astrophysics.

5.8.1 Summary

5.8.1.1 The physics of streams formation

In §5.2, we studied the formation of streams by considering the motion of small clusters of test particles, when described by action-angle coordinates of the host potential. We recalled from §8.3.1 of Binney & Tremaine (2008) that streams form in angle-space, when an initial, small structure in action-space results in a corresponding structure in frequency-space. This structure in frequency-space gives rise to a structure in angle space, which grows secularly and quickly dominates the original angle-space form of the cluster. Since the particles only feel the gravity of the host, the action-space structure remains frozen.

We show that the frequency-space structure is given by the transformation of the action-space structure under the linear map $\mathbf{D} = \nabla_{\mathbf{J}} H(\mathbf{J})$, which is the Hessian of the Hamiltonian in the actions $\mathbf{J}$. This linear map can be characterized by the image produced from a unit sphere in action-space under its transformation. The sphere will be transformed into an ellipsoid, in which directions of the semi-axes correspond to the eigenvectors of the map, and the lengths of the semi-axes correspond to the eigenvalues of the map. Streams form when one of the eigenvalues of the map—and hence, the magnitude of the stretch in the principal direction—is very much larger than the other two.

5.8.1.2 Stream formation in systems of two actions

In §5.3 we investigate the form of this linear map for the general case of a system described by two actions. We find explicit relations between the geometry of the map and the form of the Hamiltonian, when written as a function of actions. We find that the Hamiltonians of the Kepler potential and of the spherical harmonic oscillator take a form in which the geometry of the linear map is globally consistent.

In the former case, one of the eigenvalues of the map is null, so the angle-space image of any action-space structure, under this map, will always be a perfect filament. We show that this filament will always be aligned precisely with the frequency vector of the progenitor orbit. Hence, disrupting clusters always form streams in the Kepler potential, and those streams always perfectly delineate the progenitor orbit. We go on to confirm this prediction numerically.

In the case of the spherical harmonic oscillator we show that the eigenvalues of the map are both null. Hence, the image of an action-space structure under this map is trivially null. The interpretation of this result is that streams cannot form in spherical harmonic oscillator potentials; all angle-space structure are perfectly preserved.

We also consider the case of the isochrone potential, for which the Kepler potential and the spherical harmonic oscillator are limiting cases. We find that in general the eigenvalues of the map are non-zero, but that the ratio between them is everywhere large. Hence, disrupting clusters will always form streams in this potential. Unlike with the Kepler potential, we find that the principal direction of the map is not generally aligned with the frequency vector of the progenitor orbit. Hence, the angle-space stream formed from a spherical cluster in action-space will also not be aligned with that frequency vector. We find that the misalignment in angle-space will typically be of order a few degrees. We go on to confirm this prediction numerically. We also show that this angle-space misalignment manifests itself in real-space as a failure of the progenitor orbit to delineate the track of the stream.

5.8.1.3 Considerations for the mapping of streams from action-angle space into real-space

In §5.4 we deduce general limits on the anisotropy of the action-space structure formed from a disrupting cluster. In doing so, we address the concern that a stream of arbitrary angle-space orientation could be created from a sufficiently non-isotropic action-space structure. We find that any realistic cluster, if it forms a stream at all, will form a stream that is not misaligned from the frequency vector, in angle space, by more than a few degrees.

We then take this result and examine the real-space tracks of lines in angle-space that are misaligned from the frequency vector by a few degrees. We find that the nature of the anomaly between the real-space track of such a line, and the orbit to which the frequency vector corresponds, depends on the phase of the orbit at which the observation is made. Specifically, we find that if the observation is made with the orbit close to apsis, then the angle-space misalignment manifests itself as a change in curvature of the stream track, with respect to the orbit. We find that if the observation is made at a point far from apsis, then the angle-space misalignment corresponds to a real-space misalignment, of similar magnitude.

The real-space position of a star is a function of both its action and its angle. Since it is a finite structure in action-space that gives rise to an angle-space stream, it follows that the action of the stars in a stream must vary down the track. In order to predict the real-space track of a stream, it is not then sufficient to know only the angle-space structure; one must also account for this variation in action along the angle-space stream. We investigate this variation, and deduce a method to detect when it will be of consequence for predicting the real-space track of the stream. We also show that the variation can be predicted, provided one can guess the time since the cluster's first tidal stripping event, and provided one has an appropriate model of the action-space structure of the stream.

5.8.1.4 The consequences of fitting orbits using misaligned streams

In §5.5, we examine the consequences of attempting to utilize a misaligned stream to constrain the parameters of an isochrone potential, while assuming that the track does properly delineate an orbit.

When the stream perfectly delineates its progenitor orbit, we find that the potential parameters are isolated perfectly. Thus, we demonstrate the efficacy of our optimization technique. When the real-space curvature of the stream track is less than that of the progenitor orbit, we find that the mass parameter of the potential is consistently under-reported, by 11 per cent in our example case. When the curvature of the stream track is greater than that of the progenitor orbit, we find that the mass parameter is consistently over-reported, by 14 per cent in our example case.

However, in deducing these results, we have allowed ourselves to consider unrealistic combinations of parameters, which would be ruled out by independent observations. We therefore repeat the exercise, but

we now consider only that family of potential parameters which correctly reproduces a fiducial circular velocity at the Solar radius.

When the stream perfectly delineates the progenitor orbit, we again find that the potential parameters are correctly identified. However, when the stream track has less curvature than the orbit, the best-fitting potential has a mass that is 21 per cent smaller than the truth. When the stream track has greater curvature than the orbit, the best-fitting potential has a mass parameter that is 54 per cent larger than the truth. Hence, adding extra constraints to the fitting process has resulted in the systematic errors becoming worse.

In summary, we find that large systematic errors can be made when attempting to optimize potential parameter by assuming that streams act as proxies for orbits, when they do not.

5.8.1.5 The action-space distribution of N-body clusters

In §5.6 we utilized N-body simulation of King model clusters to examine the action-space distribution resulting from the disruption of a live cluster.

We found that the action-space distribution of a disrupted cluster takes a characteristic shape, which is flattened, and is oriented with a small positive gradient when viewed in in the (L, J_r) plane. We find that the greater the eccentricity of the cluster orbit, the larger this positive gradient will be, but that in general, the distribution will be roughly oriented in the $\hat{L}$ direction in action-space. We are able to explain all features of this distribution in terms of the basic physical processes that apply to clusters undergoing tidal disruption. Hence, we predict that the action-space distribution of all disrupting clusters will take the same basic form, although details such as the density of stars across the distribution, and their precise dimensions and orientation, necessarily depend on the details of the model, the orbit, and the potential in question.

We show by simulation that real disrupted clusters do indeed form streams with tracks that are poorly predicted by the progenitor orbit. However, we show that by utilizing a simple, straight-line model of the action-space distribution, we are able to predict the real-space stream track of a stream with perfect accuracy.

5.8.1.6 Non-spherical systems

In §5.7, we extend the results of the previous sections to systems with a Hamiltonian described by three actions, allowing us to consider the stream-forming properties of non-spherical systems. As our example, we utilize oblate, axisymmetric Stäckel models with asymptotically logarithmic density profiles (de Zeeuw et al., 1986). We consider two such models: one highly flattened, and therefore representative of a heavy

galactic disk; and one that is only slightly flattened, but is roughly consistent with the observed Milky Way rotation curve from the Solar circle outwards.

In both potentials, we find that the ratio of the eigenvalues of the map is very large; hence, disrupted clusters will always form streams in these potentials. We further find that the principal direction of the map is misaligned with the frequency vector by $\sim 1^\circ$ when the flattening is only slight, and by $\sim 10^\circ$ when the flattening is substantial. On this basis, we generally expect streams not to be well represented by orbits in these potentials, and we expect the representation to be worse when the potential is flatter.

We performed an N-body simulation of a disrupting cluster on an approximately planar orbit in the highly flattened Stäckel potential. We found that the angle-space misalignment between the stream and the frequency vector is indeed large, and that the resulting real-space track is very poorly represented by the progenitor orbit. However, we again find that a simple straight-line model of the action-space distribution predicts the corresponding real-space stream track with superb accuracy.

We took the less flattened Stäckel potential to be a model of the Milky Way potential, and in this potential we performed N-body simulations of streams that are superficially similar to the observed Milky Way streams GD-1 (Grillmair & Dionatos, 2006) and the Orphan stream (Belokurov et al., 2007). We found that, in both cases, the angle-space stream and the frequency vector were almost perfectly aligned. This occurs because of a property of spherical logarithmic potentials, which causes the action-space image of the frequency vector to align closely with the action-space distribution of a disrupted cluster.

In the case of GD-1, this results in a stream that is almost perfectly represented by its progenitor orbit. Hence, we found that the assumption made in Chapter 4 that the GD-1 stream perfectly delineates the orbits of its stars was a fortuitously fair one, and that the analysis of the Galactic parallax of GD-1 therefore needs no revision in light of this work. However, this conclusion comes with the caveat that the Stäckel model used is not a particularly realistic model of the Milky Way potential, and the result should be confirmed with a better model.

In the case of the Orphan stream, we find that despite the stream and the frequency vector being perfectly aligned in angle-space, the real-space track of the stream is still not well represented by the trajectory of the progenitor orbit. This occurs because the variation in action down the stream is sufficient to alter the curvature of the stream track near apsis, resulting in the divergence of the stream track and the orbit. However, we find once again that a simple model of the action-space distribution allows us to predict this stream track with perfect accuracy.

Lastly, we found that in both our Stäckel models, the action-space distribution resulting from the disruption of clusters is directly comparable to that found in the isochrone potential, confirming the generality of those observations.

5.8.2 Discussion and future directions

We have found that stream tracks cannot be used as reliable proxies for orbits without detailed consideration of the mechanics of their formation. One immediate consequence is to generally render as unreliable those techniques that attempt to constrain the parameters of the Galactic potential (e.g. Newberg et al., 2010, Willett et al., 2009 and Koposov et al., 2010) without an appropriate analysis of whether the stream is actually well represented by its progenitor orbit.

Such a one-off analysis could be performed using N-body simulation, although utilizing the techniques presented in this chapter would be quicker. If the simulation confirms that the stream is well modelled by an orbit, then one may proceed as before. However, if the stream is not well modelled by an orbit, as will generally be the case, the technique of fitting orbits can no longer be used.

In such circumstances, one might resort to N-body shooting methods, such as Johnston et al. (2005), to compute stream tracks to feed to an optimization algorithm. However, the problem with such methods is the sensitivity of the output to the detail of the initial conditions. One must perform a multitude of simulations, over a range of initial conditions, in order to reject a single potential from further consideration. This requires both great computational resources, and even then, it is impossible to consider more than a tiny fraction of the $\sim 10^7$ candidate orbits that we found it necessary to consider in order to fully test the parameter space and definitively exclude a potential (Chapter 2). The conclusions drawn from such as an exercise could therefore only be weak, and therefore unsatisfactory.

The results of this chapter present a possible alternative. We have found that with simple models of the action-space distribution of a disrupted cluster, such as can be readily obtained from a single N-body simulation, we can reliably and accurately predict the track of a stream, even when it diverges significantly from the trajectory of the progenitor orbit. In principle, we can compute these tracks with no more computational effort than it takes to integrate an orbit. Hence, the techniques for fitting orbits, such as those presented in the earlier chapters of this thesis, could be readily adapted to fit stream tracks instead.

To achieve this goal, the following hurdles need to be overcome:

1. We must achieve a detailed understanding of the action-space structure of a disrupted cluster, for any problem parameters of our choice.

   One approach would be to use N-body simulations to obtain the action-space distribution for a small number of cluster models on a set of possible orbits, in a given potential. The resulting distributions could be used as a basis set, which would be interpolated and distorted to provide an estimate for the action-space structure for any given cluster on any chosen orbit. The required distortions for changes to cluster model and orbit parameters have already been touched upon in

this chapter, although for general applicability, a complete quantitative theory of these distortions will be required.

However achieved, a rigorous and systematic schema for efficiently modelling the action-space structure of clusters must be developed if our insights are to be practically useful.

2. Application of the techniques of this chapter requires the ability to compute action-angle variables from conventional phase-space coordinates, and vice versa, with reasonable accuracy. In this chapter, we have restricted ourselves to those few potentials in which the transformation can be readily made. In general, one would like to work with more sophisticated potentials, for which no easy transformation between action-angle variables and phase-space coordinates can be made.

   Fortunately, the "torus machine" of McMillan & Binney (2008), which builds on the prior work of Kaasalainen & Binney (1994) and McGill & Binney (1990), enables the actions, the frequencies and their derivatives to be accurately and quickly computed for regular orbits in realistic Galactic potentials. Combined with the torus machine, the techniques explored in this chapter could be extended to work with such models.

3. We have seen that, in certain circumstances, the precise trajectory of a stream can be sensitive to the variation in the actions of the stars along the stream. When this is the case, it is necessary to know the elapsed period since the first pericentre passage of a cluster, in order to predict the resulting stream track with maximum precision. Further work is necessary to discover to what extent it is possible to estimate this elapsed period from observations of resulting streams themselves.

   Unfortunately, it is likely that when these certain circumstances prevail, efforts to correctly predict the entirety of the stream track will prove to be error prone. However, it should always be possible to accurately predict those parts of the track that are most likely to be erroneous. Work should therefore be done to provide a quick and robust method for detecting when these certain circumstances arise, and which parts of the predicted track are affected, such that appropriate levels of uncertainty can be attached to the erroneous regions.

4. Given success in items 1–3 above, we will be able to compute the track of a stream, given an initial condition, with similar computational expenditure as is currently required to compute an orbital trajectory from the same.

   It would then require only a small technical change to alter the many techniques that attempt to fit orbits to observations of streams, for instance those of Willett et al. (2009), in order to have them fit accurate stream tracks instead.

Many of these techniques already have the power to diagnose the Galactic potential, based on their fitting of orbits. Once converted to fit stream tracks instead, the existing code could be utilized directly to constrain the Galactic potential, with neither the assumption that streams follow orbits, nor the accompanying risk of erring should they not.

A suggested programme of immediate further work is to tackle items 1–3 from this list in sequence. The resulting stream-predicting engine can then be coupled to existing orbit-fitting techniques, and immediately applied to the Orphan stream data of Newberg et al. (2010) and the GD-1 data of Koposov et al. (2010).

In the future, the techniques presented here may well be applicable to the Sagittarius dwarf stream. In this chapter we have assumed that our low-mass clusters do not affect the host potential. This may not be true in the case of a heavy Sagittarius progenitor. Further study of the effect of a live host potential on the mechanics of stream formation will be required for the techniques to be reliably applicable to the Sagittarius dwarf stream.

Chapter 6

Conclusions

6.1 Overview

The automated Sloan Digital Sky Survey (SDSS, York et al., 2000) has revealed a tremendous amount of substructure in the stellar halo of the Milky Way Galaxy. When appropriate cuts to the data are made, the halo is revealed to be streaked with streams formed from the shattered remains of galaxies, once captured by the Milky Way and now in the process of being subsumed (Fig. 1.1; Belokurov et al., 2006). Careful analysis of the data from SDSS and its follow-on extension programme (SEGUE, Yanny et al., 2009) has uncovered large numbers of these streams (Odenkirchen et al., 2003; Majewski et al., 2003; Yanny et al., 2003; Belokurov et al., 2006, 2007; Grillmair, 2006a; Grillmair & Dionatos, 2006; Grillmair & Johnson, 2006; Grillmair, 2009; Newberg et al., 2009).

Tidal streams have been empirically noted to delineate the orbits of their progenitors (McGlynn, 1990; Johnston et al., 2005). Given this conjecture, the diagnostic power that a single stream provides over the form of the Galactic potential is extraordinary (Binney, 2008). Unfortunately, simulation results have latterly revealed that this conjecture is untrue: tidal streams do not delineate individual orbits (Choi et al., 2007; Chapter 2).

The subject of this thesis has been the study of these streams, and how we can best unlock their diagnostic power, given real data from streams that may not precisely delineate an orbit. In pursuit of this goal, we have examined the following problems.

6.1.1 Fitting orbits using radial velocity data

The conventional approach to exploit the diagnostic power of streams is as follows: guess a trial potential, and then integrate orbits from a range of initial conditions in order to identify those orbits that best

match the data. In practice, this technique is difficult to implement, because the combined space of initial conditions and trial potentials is too vast to search effectively. The search rarely yields an orbits that is a satisfactory fit to the data, and even if it does, there is no guarantee that the correct orbit/potential combination has been chosen.

One alternative are so-called geometrodynamical techniques, which utilize measured velocity data in order to reduce the scope of the orbit-fitting problem (Jin & Lynden-Bell, 2007; Binney, 2008; Jin & Lynden-Bell, 2008). In Chapter 2 we advance the work of Binney (2008) and Jin & Lynden-Bell (2007) by adapting the geometrodynamical reconstruction algorithm they present, such that it is able to handle erroneous input data. It is not necessary to specify the nature of the errors in order to proceed: they can be systematic or statistical or a combination of both. In particular, it is no longer necessary to assume that streams perfectly delineate orbits. However, in all cases the errors must be contained within a specified boundary.

We find that we can successfully overcome the limitations the prior methods face when handling real data, and that we can identify those orbits—should any exist—that are consistent with a given potential and the input data. We further find that, given sufficiently precise input data, the ability shown by the Binney (2008) algorithm to diagnose the correct form of the potential is retained.

6.1.2 Fitting orbits using proper motion data

The key limitation on the work in Chapter 2 is the lack of availability of line-of-sight velocity data for the hundreds of main-sequence stars that make up a stream. It is unlikely that sufficient 8-m class telescope time will be afforded, in the near future, to remedy this deficit for more than a few example cases. Conversely, the Pan-STARRS survey (Kaiser et al., 2002, currently commissioning), the Gaia mission (Perryman et al., 2001) and the LSST (Tyson, 2002) will soon produce catalogues of proper motions for billions of main-sequence stars in the Milky Way.

In Chapter 3, we complement the work of Binney (2008), by creating a geometrodynamical algorithm to reconstruct orbits by utilizing proper-motion measurements instead of radial-velocity measurements. We showed that this algorithm retains the ability of the radial-velocity algorithm to diagnose the Galactic potential.

6.1.3 Galactic parallax

In Chapter 4 we explore a technique, which arises out of the work of Chapter 3, that permits the measurement of distances to remote stars in streams given only the measured proper motions of those stars, and no assumption being made about the form of the potential.

The technique utilizes an effect that we call Galactic parallax: the apparent motion of stars in a direction other than along their stream due to the reflex motion of the Sun. Given knowledge of the velocity of the Sun with respect to the Galactic centre, this effect enables trigonometric distances to be calculated for stars far beyond the range of conventional parallax.

We have examined in detail the practicality of this technique, and we demonstrate its use by measuring the distance to the tidal stream GD-1.

6.1.4 The mechanics of streams

In Chapter 5, we studied the mechanics of stream formation from first principles. Our motivation for doing so was to investigate under what circumstances streams could be relied upon to delineate orbits.

We discovered that, in general, streams do not precisely delineate orbits. The degree to which they do not is dependent upon a number of factors: the shape of the potential, the orbit of the stream's progenitor, and the size and shape of the progenitor itself.

We find that the real-space manifestation of this failure depends upon the phase at which the stream is observed. Specifically, if the stream is observed away from apsis, the stream will be misaligned with its orbit, while if the stream is observed close to apsis, the stream will display a different curvature to the orbit.

We find that constraining the parameters of a the Galactic potential by using a stream—while wrongly assuming that the stream delineates an orbit—can cause large systematic errors in the reported parameters.

However, we do find that, given a simple model of the phase-space distribution of a stream that is informed by the results of this chapter, we are able to predict the real-space tracks of streams with high accuracy, even when those streams are poorly represented by an orbit.

6.2 This work in context

6.2.1 The dark matter distribution of the Galaxy

The dark matter distribution in the Milky Way Galaxy is still very much unknown, and devising techniques to effectively probe it is a key challenge in galactic astrophysics. Attempts to directly detect the annihilation of dark matter particles have failed thus far, meaning that the only probe for the distribution of dark matter in our Galaxy is its gravitational effect on the dynamics of baryonic matter.

The recently discovered wealth of substructure in the form of "fossil relics" from merger events in the Galactic environment provides an effective probe of the gravitational field of the Galaxy, and therefore

of the dark matter distribution as well. Many standard techniques exist that attempt to harness streams to probe the potential, and almost all involve fitting orbits to the streams.

The principles of stream formation that we have elucidated in this thesis should now raise caution in any attempt to utilize streams as environmental probes. In particular, the failure of streams to precisely delineate orbits can cause serious systematic error when attempting to constrain potential parameters. It may be possible, with further work, to repair these standard techniques to fit stream-tracks, instead of orbits, to their target data. Otherwise, the use of the standard techniques will have to be restricted to those streams which can be demonstrably shown to be well represented by orbits.

Meanwhile, the alternative reconstruction techniques that we have developed may well prove key to constraining the dark matter distribution. In particular, they can compensate for the failure of streams to delineate orbits, and they have the power to show that a given potential is completely inconsistent with a stream, which would rule it out as a possible Galactic potential.

The only hurdle to the widespread application of these techniques is the lack of appropriate velocity data for the remote main-sequence stars that make up tidal streams. The shortage of radial-velocity data for such stars in the Milky Way is unlikely to be remedied soon. However, there is likely to be an explosion in the availability of high-quality proper-motion data in the near future, as several forthcoming astrometric projects come online.

Methods such as those presented in this thesis should then be immediately applied to all known tidal streams for which the data become available. This list would be quite long: GD-1, Pal 5, the Orphan stream, the tails of the cluster NGC 5466 and the Sagittarius Dwarf stream. Undoubtedly, more streams will be discovered because of the arrival of the data itself, so this list is expected to grow.

The key to rapid success will be to work out, in the time before these data arrive, how to best constrain the Galactic potential given the combination of all possible sources of information: proper-motion fitting of these streams, radial-velocity fitting of those streams for which the data is to hand, and constraints on the Galactic potential from other observations.

This is an important challenge to which Galactic astrophysics should commit itself over the next few years. If it does so, the scientific reward could well be enormous.

6.2.2 Distance estimation

Distance estimation in our Galaxy is a difficult problem. Trigonometric parallax produces accurate distance measurements that can be used to calibrate other methods, but its range is the most limited of all standard techniques. Meanwhile, distance estimation based on the analysis of starlight has practically indefinite range, but is subject to the effects of extinction and reddening caused by intervening matter,

and also to some of the unknown effects of chemical make-up on stellar luminosity.

The technique of Galactic parallax, which we present, produces distance estimates to remote stream stars that are as fundamental as conventional trigonometric parallax, but with vastly superior range. The technique is restricted, in that the effect can only be measured for stars that are part of a stream. However, the technique can be used to calibrate or independently check the calibration of other techniques, such as photometric distance estimation, which have more widespread applicability.

When high-quality proper motion data from next-generation astrometric projects becomes available in the next few years, this widespread application of this technique will be possible: the data from Gaia and Pan-STARRS will put much of the Galaxy in the range of trigonometric distance calculation for the first time.

The only caveat to its use is the requirement to predict the rest-frame direction of motion of a star from the track of its stream on the sky. In those cases where streams are well represented by orbits, estimating this direction is trivial: the star moves tangent to the stream. In the general case, where streams are not well represented by orbits, Galactic parallax can still be computed, but the techniques expounded in Chapter 5 of this thesis will need to be used to correct the for the misalignment between the motion of the star and the tangent to the stream.

6.3 Future work

Each chapter of this thesis contains its own discussion of the immediate directions for extending the work within it. We include a short summary of those possibilities below, in addition to some general observations.

1. The radial-velocity algorithm of Chapter 2 should be immediately applied to those streams for which we have data, namely, the Orphan stream (Newberg et al., 2010) and the GD-1 stream (Koposov et al., 2010). In addition, work should be undertaken to improve the efficiency of the optimization routine. Future work might also examine the routines for solving the reconstruction equations, in order to see if the noise floor affecting the calculation can be decreased further.

2. The proper-motion algorithm of Chapter 3 should be adapted to the modification-and-search routines presented in Chapter 2, so that it can then be used with real data. The resulting procedure should then be applied to the proper-motion data for the GD-1 stream (Koposov et al., 2010), with the results compared to the radial-velocity analysis of the same.

3. The technique of Galactic parallax should be immediately applied to any stream for which sufficiently accurate proper-motion measurements can be found. Further, the technique is actually

applicable to any galactic structure for which a proper motion can be measured, and which enables the rest-frame trajectory of its constituents to be predicted. It may be that some features of the Galaxy, such as spiral arms, allow the rest-frame motion of the stars to be predicted statistically: if a proper motion can also be measured for these features, then a Galactic parallax can also be computed.

4. It is important that the work in Chapter 5 on the mechanics of streams be quickly repeated in a realistic model for the Galaxy potential, so that we may assess the degree to which Milky Way streams are well-represented by orbits. This requires marriage to those techniques that allow action-angle variables to be computed in general potentials (McMillan & Binney, 2008). If we then discover that Milky Way streams do significantly deviate from orbits, attention should be paid to repairing existing potential-optimization routines by having them fit stream-tracks instead of orbits.

In 2012, the Gaia mission is expected to launch, and will return full phase-space data for over one billion Milky Way stars. Given the number of streams identified in the comparatively poor SDSS data, we may expect many more discoveries to be made from Gaia data, and many of the already-known streams will also become fully characterized.

The Gaia era will represent a golden age in Galactic astronomy, and will allow a Galactic model of unprecedented utility to be constructed, once the right tools for the task are available. In order to achieve this goal without delay, in the few years until this data-set becomes available, a substantial theoretical effort should be directed towards understanding how to combine the output of the methods presented here—and others—in a way that places the Galactic potential under the strongest possible constraints.

References

Ahmed, Z. 2009, arXiv:0912.3592

Aumer, M., & Binney, J. J. 2009, Mon. Not. R. Astron. Soc., 397, 1286

Belokurov, V., et al. 2006, Astrophys. J. Lett., 642, L137

—. 2007, Astrophys. J., 658, 337

Bethe, H. 1972, Energy Production in Stars, Nobel Lectures, Physics 1963-1970 (Amsterdam: Elsevier)

Binney, J. 2008, Mon. Not. R. Astron. Soc., 386, L47

Binney, J., & Merrifield, M. 1998, Galactic Astronomy (Princeton: Princeton University Press)

Binney, J., & Tremaine, S. 2008, Galactic Dynamics: Second Edition (Princeton: Princeton University Press)

Bullock, J. S., & Johnston, K. V. 2005, Astrophys. J., 635, 931

Bullock, J. S., Kolatt, T. S., Sigad, Y., Somerville, R. S., Kravtsov, A. V., Klypin, A. A., Primack, J. R., & Dekel, A. 2001, Mon. Not. R. Astron. Soc., 321, 559

Bullock, J. S., Kravtsov, A. V., & Weinberg, D. H. 2000, Astrophys. J., 539, 517

Choi, J., Weinberg, M. D., & Katz, N. 2007, Mon. Not. R. Astron. Soc., 381, 987

Cox, T. J., & Loeb, A. 2008, Mon. Not. R. Astron. Soc., 386, 461

de Zeeuw, T. 1985, Mon. Not. R. Astron. Soc., 216, 273

de Zeeuw, T., Peletier, R., & Franx, M. 1986, Mon. Not. R. Astron. Soc., 221, 1001

Dehnen, W., & Binney, J. 1998a, Mon. Not. R. Astron. Soc., 294, 429

Dehnen, W., & Binney, J. J. 1998b, Mon. Not. R. Astron. Soc., 298, 387

Dehnen, W., Odenkirchen, M., Grebel, E. K., & Rix, H. 2004, Astron. J., 127, 2753

Drimmel, R., & Spergel, D. N. 2001, Astrophys. J., 556, 181

Dubinski, J., Mihos, J. C., & Hernquist, L. 1996, Astrophys. J., 462, 576

Einstein, A. 1915, Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin), 844

Eyre, A. 2010, Mon. Not. R. Astron. Soc., 403, 1999

Eyre, A., & Binney, J. 2009a, Mon. Not. R. Astron. Soc., 399, L160

—. 2009b, Mon. Not. R. Astron. Soc., 400, 548

Faber, S. M., & Gallagher, J. S. 1979, Ann. Rev. Astron. Astrophys., 17, 135

Fellhauer, M., Evans, N. W., Belokurov, V., Wilkinson, M. I., & Gilmore, G. 2007a, Mon. Not. R. Astron. Soc., 380, 749

Fellhauer, M., et al. 2007b, Mon. Not. R. Astron. Soc., 375, 1171

Gamow, G. 1970, My World Line: An Informal Autobiography (New York: Viking Press)

Ghez, A. M., Salim, S., Hornstein, S. D., Tanner, A., Lu, J. R., Morris, M., Becklin, E. E., & Duchêne, G. 2005, Astrophys. J., 620, 744

Gillessen, S., Eisenhauer, F., Trippe, S., Alexander, T., Genzel, R., Martins, F., & Ott, T. 2009, Astrophys. J., 692, 1075

Grillmair, C. J. 2006a, Astrophys. J. Lett., 645, L37

—. 2006b, Astrophys. J. Lett., 651, L29

—. 2009, Astrophys. J., 693, 1118

Grillmair, C. J., Carlin, J. L., & Majewski, S. R. 2008, Astrophys. J. Lett., 689, L117

Grillmair, C. J., & Dionatos, O. 2006, Astrophys. J. Lett., 643, L17

Grillmair, C. J., & Johnson, R. 2006, Astrophys. J. Lett., 639, L17

Helmi, A. 2008, Astron. Astrophys. Rev., 15, 145

Helmi, A., & White, S. D. M. 1999, Mon. Not. R. Astron. Soc., 307, 495

Helmi, A., White, S. D. M., de Zeeuw, P. T., & Zhao, H. 1999, Nature, 402, 53

Hubble, E. 1929, Proceedings of the National Academy of Science, 15, 168

Hubble, E. P. 1926, Astrophys. J., 64, 321

Ibata, R. A., Gilmore, G., & Irwin, M. J. 1995, Mon. Not. R. Astron. Soc., 277, 781

Ivezić, Ž., et al. 2008, in IAU Symposium, ed. W. J. Jin, I. Platais, & M. A. C. Perryman, Vol. 248, 537–543

Jin, S., & Lynden-Bell, D. 2007, Mon. Not. R. Astron. Soc., 378, L64

—. 2008, Mon. Not. R. Astron. Soc., 383, 1686

Johnston, K. V., Hernquist, L., & Bolte, M. 1996, Astrophys. J., 465, 278

Johnston, K. V., Law, D. R., & Majewski, S. R. 2005, Astrophys. J., 619, 800

Johnston, K. V., Spergel, D. N., & Hernquist, L. 1995, Astrophys. J., 451, 598

Jurić, M., et al. 2008, Astrophys. J., 673, 864

Kaasalainen, M., & Binney, J. 1994, Mon. Not. R. Astron. Soc., 268, 1033

Kaiser, N., et al. 2002, in SPIE Conference Series, ed. J. A. Tyson & S. Wolff, Vol. 4836, 154–164

King, I. R. 1966, Astron. J., 71, 64

Klypin, A., Kravtsov, A. V., Valenzuela, O., & Prada, F. 1999, Astrophys. J., 522, 82

Koposov, S. E., Rix, H., & Hogg, D. W. 2010, Astrophys. J., 712, 260

Laughlin, G. 2009, Nature, 459, 781

Law, D. R., Johnston, K. V., & Majewski, S. R. 2005, Astrophys. J., 619, 807

Londrillo, P., Nipoti, C., & Ciotti, L. 2003, Memorie della Societa Astronomica Italiana Supplement, 1, 18

Magnier, E. A., Liu, M., Monet, D. G., & Chambers, K. C. 2008, in IAU Symposium, ed. W. J. Jin, I. Platais, & M. A. C. Perryman, Vol. 248, 553–559

Majewski, S. R., Skrutskie, M. F., Weinberg, M. D., & Ostheimer, J. C. 2003, Astrophys. J., 599, 1082

Mathewson, D. S., Cleary, M. N., & Murray, J. D. 1974, Astrophys. J., 190, 291

McGill, C., & Binney, J. 1990, Mon. Not. R. Astron. Soc., 244, 634

McGlynn, T. A. 1990, Astrophys. J., 348, 515

McMillan, P. J., & Binney, J. J. 2008, Mon. Not. R. Astron. Soc., 390, 429

—. 2010, Mon. Not. R. Astron. Soc., 402, 934

Miyamoto, M., & Nagai, R. 1975, Publ. Aston. Soc. Jpn., 27, 533

Montuori, M., Capuzzo-Dolcetta, R., Di Matteo, P., Lepinette, A., & Miocchi, P. 2007, Astrophys. J., 659, 1212

Moore, B., & Davis, M. 1994, Mon. Not. R. Astron. Soc., 270, 209

Moore, B., Diemand, J., Madau, P., Zemp, M., & Stadel, J. 2006, Mon. Not. R. Astron. Soc., 368, 563

Moore, B., Governato, F., Quinn, T., Stadel, J., & Lake, G. 1998, Astrophys. J. Lett., 499, L5+

Munn, J. A., et al. 2004, Astron. J., 127, 3034

Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, Astrophys. J., 490, 493

Newberg, H. J., Willett, B. A., Yanny, B., & Xu, Y. 2010, Astrophys. J., 711, 32

Newberg, H. J., Yanny, B., & Willett, B. A. 2009, Astrophys. J. Lett., 700, L61

Odenkirchen, M., Grebel, E. K., Kayser, A., Rix, H., & Dehnen, W. 2009, Astron. J., 137, 3378

Odenkirchen, M., et al. 2003, Astron. J., 126, 2385

Ostriker, J. P., & Hausman, M. A. 1977, Astrophys. J. Lett., 217, L125

Perryman, M. A. C., et al. 2001, Astron. Astrophys., 369, 339

Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 2002, Numerical Recipes in C++. The Art of Scientific Computing: Second Edition (Cambridge: Cambridge University Press)

Reid, M. J., & Brunthaler, A. 2004, Astrophys. J., 616, 872

Rix, H., de Zeeuw, P. T., Cretton, N., van der Marel, R. P., & Carollo, C. M. 1997, Astrophys. J., 488, 702

Sales, L. V., et al. 2008, Mon. Not. R. Astron. Soc., 389, 1391

Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, Astrophys. J., 500, 525

Skrutskie, M. F., et al. 2006, Astron. J., 131, 1163

Smith, M. C., et al. 2007, Mon. Not. R. Astron. Soc., 379, 755

Soares, D. S. L. 2001, J. R. Astron. Soc. Can., 95, 10

Spitzer, L. 1987, Dynamical Evolution of Globular Clusters (Princeton: Princeton University Press)

Toomre, A., & Toomre, J. 1972, Astrophys. J., 178, 623

Tremaine, S. 1999, Mon. Not. R. Astron. Soc., 307, 877

Tyson, J. A. 2002, in SPIE Conference Series, ed. J. A. Tyson & S. Wolff, Vol. 4836, 10–20

van Leeuwen, F. 2007, Hipparcos, the New Reduction of the Raw Data, Astrophysics and Space Science Library (Dordrecht: Springer)

Velazquez, H., & White, S. D. M. 1995, Mon. Not. R. Astron. Soc., 275, L23

Vergely, J., Ferrero, R. F., Egret, D., & Koeppen, J. 1998, Astron. Astrophys., 340, 543

White, S. D. M., & Rees, M. J. 1978, Mon. Not. R. Astron. Soc., 183, 341

Willett, B. A., Newberg, H. J., Zhang, H., Yanny, B., & Beers, T. C. 2009, Astrophys. J., 697, 207

Wright, T. 1750, An Original Theory or New Hypothesis of the Universe (London)

Yanny, B., et al. 2003, Astrophys. J., 588, 824

—. 2009, Astron. J., 137, 4377

York, D. G., et al. 2000, Astron. J., 120, 1579

Zwicky, F. 1933, Helvetica Physica Acta, 6, 110

Appendix A

Some results in Stäckel potentials

This appendix contains some useful results in Stäckel potential calculations. They are based on the work of de Zeeuw (1985, herein Z85), but to the best of our knowledge, most of them have not been published, and so we include them here. Those that have been published but are included nonetheless are required in order to inform the later results. In order to maintain consistency, we adopt the notation from Z85. Furthermore, to remain concise, we will not explain this notation, except where it departs from that presented in Z85. Reference should be made to that work at all times.

A.1 Computing the actions

Equation (118) from Z85 shows that,

$$J_\tau = \frac{1}{2\pi} \oint p_\tau \, \mathrm{d}\tau. \tag{A.1}$$

In asking a computer to perform this quadrature, we want to make the integrand as flat as possible, with respect to the quadrature scheme that we are pursuing. In this case, we only expect ill-conditioned behaviour near the endpoints, (τ_0, τ_1). We re-write equation (A.1)

$$J_\tau = \frac{1}{2\pi} \oint \sqrt{(\tau - \tau_0)(\tau_1 - \tau)} \tilde{p}(\tau) \, \mathrm{d}\tau, \tag{A.2}$$

where $\tilde{p}$ is a well-behaved function of τ that becomes flat near the endpoints (τ_0, τ_1).

Although the integrand of equation (A.2) is well-behaved everywhere, we would like to remove explicit dependence upon radicals of the dummy variable. This is because any numerical noise near the endpoints could send the argument of these radicals negative, which would not compute.

We write

$$\bar{\tau} = \frac{1}{2}(\tau_0 + \tau_1); \qquad \hat{\tau} = \frac{1}{2}(\tau_1 - \tau_0); \qquad \tau = \hat{\tau} \sin\theta + \bar{\tau}. \tag{A.3}$$

Equation (A.2) becomes,

$$J_\tau = \frac{1}{2\pi} \oint \hat{\tau}^2 \cos^2\theta \, \tilde{p}(\tau(\theta)) \, \mathrm{d}\theta, \tag{A.4}$$

where we observe that

$$(\tau - \tau_0)(\tau_1 - \tau) = \hat{\tau}^2 \cos^2\theta, \tag{A.5}$$

and the integral, which was over $\tau = (\tau_0, \tau_1)$, is now over $\theta = (-\pi/2, \pi/2)$.

A.2 Computing the frequencies

Equations 125 and 126 in Z85 show that the frequencies are given by,

$$\Omega_\lambda = \frac{1}{\Delta}\frac{\partial(J_\mu, J_\nu)}{\partial(I_2, I_3)}, \qquad \Omega_\mu = \frac{1}{\Delta}\frac{\partial(J_\nu, J_\lambda)}{\partial(I_2, I_3)}, \qquad \Omega_\nu = \frac{1}{\Delta}\frac{\partial(J_\lambda, J_\mu)}{\partial(I_2, I_3)}, \tag{A.6}$$

where we have defined the determinant,

$$\Delta = \frac{\partial(J_\lambda, J_\mu, J_\nu)}{\partial(E, I_2, I_3)}. \tag{A.7}$$

We must calculate explicitly the partial derivatives of the actions J_τ wrt to the integrals (E, I_2, I_3) by explicitly differentiating the equations (A.1). We find,

$$\frac{\partial J_\tau}{\partial E} = \frac{1}{8\pi}\oint \frac{\mathrm{d}\tau}{(\tau+\beta)p_\tau}, \tag{A.8}$$

$$\frac{\partial J_\tau}{\partial I_2} = -\frac{1}{8\pi}\oint \frac{\mathrm{d}\tau}{(\tau+\alpha)(\tau+\beta)p_\tau}, \tag{A.9}$$

$$\frac{\partial J_\tau}{\partial I_3} = -\frac{1}{8\pi}\oint \frac{\mathrm{d}\tau}{(\tau+\gamma)(\tau+\beta)p_\tau}, \tag{A.10}$$

where we have noted that

$$\begin{aligned}
\frac{\mathrm{d}p_\tau}{\mathrm{d}E} &= \frac{1}{4p_\tau(\tau+\beta)}, \\
\frac{\mathrm{d}p_\tau}{\mathrm{d}I_2} &= -\frac{1}{4p_\tau(\tau+\beta)(\tau+\alpha)}, \\
\frac{\mathrm{d}p_\tau}{\mathrm{d}I_3} &= -\frac{1}{4p_\tau(\tau+\beta)(\tau+\gamma)}.
\end{aligned} \tag{A.11}$$

The factors of $(\tau+\alpha)$, etc, in the denominators of these expressions are well-behaved over the entire range of integration $\tau = (\tau_0, \tau_1)$. The factors of $1/p_\tau$, however, give integrable singularities at the ends of the range, since $p_\tau^2 \propto (\tau - \tau_0)$ as $\tau \to \tau_0$.

We approach the problem using the same coordinate transformation as above (equation A.3). Equation (A.8) becomes,

$$\frac{\partial J_\tau}{\partial E} = \frac{1}{8\pi}\oint \frac{\hat{\tau}\cos\theta \mathrm{d}\theta}{(\tau+\beta)\hat{\tau}\tilde{p_\tau}\cos\theta} = \frac{1}{8\pi}\oint \frac{\mathrm{d}\theta}{(\tau+\beta)\tilde{p_\tau}}. \tag{A.12}$$

This expression can now be accurately evaluated using standard quadrature algorithms. Similarly, equa-

tions (A.9) and (A.10) are written,

$$\frac{\partial J_\tau}{\partial I_2} = -\frac{1}{8\pi}\oint \frac{\mathrm{d}\theta}{(\tau+\alpha)(\tau+\beta)\tilde{p_\tau}}, \tag{A.13}$$

$$\frac{\partial J_\tau}{\partial I_3} = -\frac{1}{8\pi}\oint \frac{\mathrm{d}\theta}{(\tau+\gamma)(\tau+\beta)\tilde{p_\tau}}. \tag{A.14}$$

A.3 Computing the derivatives of the frequencies

For the purposes of computing the Hessian $\mathbf{D}$ in Chapter 5, it is necessary to be able to evaluate the derivatives $\partial\Omega_i/\partial J_j$.

For a function $X(E, I_2, I_3)$ of the integrals (E, I_2, I_3), the definition of the partial derivative gives

$$\begin{aligned}\frac{\partial X}{\partial J_\tau} &= \frac{\partial X}{\partial E}\frac{\partial E}{\partial J_\tau} + \frac{\partial X}{\partial I_2}\frac{\partial I_2}{\partial J_\tau} + \frac{\partial X}{\partial I_3}\frac{\partial I_3}{\partial J_\tau} \\ &= \frac{\partial X}{\partial E}\Omega_\tau + \frac{\partial X}{\partial I_2}\frac{\partial I_2}{\partial J_\tau} + \frac{\partial X}{\partial I_3}\frac{\partial I_3}{\partial J_\tau}.\end{aligned} \tag{A.15}$$

The derivatives of the frequencies wrt the actions can thus be calculated in terms of the derivatives of the frequencies wrt the integrals, and the derivatives of the integrals wrt the actions.

A.3.1 The derivatives of the integrals wrt the actions

The derivatives of the integrals wrt the actions can be computed as follows. We may write

$$\begin{aligned}\frac{\partial I_2}{\partial E} &= \frac{\partial I_2}{\partial J_\lambda}\frac{\partial J_\lambda}{\partial E} + \frac{\partial I_2}{\partial J_\mu}\frac{\partial J_\mu}{\partial E} + \frac{\partial I_3}{\partial J_\lambda}\frac{\partial J_\nu}{\partial E} = 0, \\ \frac{\partial I_2}{\partial I_2} &= \frac{\partial I_2}{\partial J_\lambda}\frac{\partial J_\lambda}{\partial I_2} + \frac{\partial I_2}{\partial J_\mu}\frac{\partial J_\mu}{\partial I_2} + \frac{\partial I_3}{\partial J_\lambda}\frac{\partial J_\nu}{\partial I_2} = 1, \\ \frac{\partial I_2}{\partial I_3} &= \frac{\partial I_2}{\partial J_\lambda}\frac{\partial J_\lambda}{\partial I_3} + \frac{\partial I_2}{\partial J_\mu}\frac{\partial J_\mu}{\partial I_3} + \frac{\partial I_3}{\partial J_\lambda}\frac{\partial J_\nu}{\partial I_3} = 0,\end{aligned} \tag{A.16}$$

The derivatives of I_2 wrt the actions are then given by the cofactor expressions,

$$\frac{\partial I_2}{\partial J_\lambda} = \frac{1}{\Delta}\frac{\partial(J_\nu, J_\mu)}{\partial(E, I_3)}; \qquad \frac{\partial I_2}{\partial J_\mu} = \frac{1}{\Delta}\frac{\partial(J_\lambda, J_\nu)}{\partial(E, I_3)}; \qquad \frac{\partial I_2}{\partial J_\lambda} = \frac{1}{\Delta}\frac{\partial(J_\mu, J_\lambda)}{\partial(E, I_3)}. \tag{A.17}$$

Similarly, the derivatives of I_3 are given by the cofactor expressions,

$$\frac{\partial I_3}{\partial J_\lambda} = \frac{1}{\Delta}\frac{\partial(J_\mu, J_\nu)}{\partial(E, I_2)}; \qquad \frac{\partial I_3}{\partial J_\mu} = \frac{1}{\Delta}\frac{\partial(J_\nu, J_\lambda)}{\partial(E, I_2)}; \qquad \frac{\partial I_3}{\partial J_\lambda} = \frac{1}{\Delta}\frac{\partial(J_\lambda, J_\mu)}{\partial(E, I_2)}. \tag{A.18}$$

A.3.2 The derivatives of the frequencies wrt the integrals

The derivatives of the frequencies wrt the integrals are computed as follows. Differentiating equation (A.6) wrt some integral Z, we find

$$\begin{aligned}\frac{\partial \Omega_\lambda}{\partial Z} =& -\frac{1}{\Delta^2}\frac{\partial \Delta}{\partial Z}\frac{\partial(J_\mu, J_\nu)}{\partial(I_2, I_3)} \\ &+ \frac{1}{\Delta}\left\{\frac{\partial J_\nu}{\partial I_3}\frac{\partial^2 J_\mu}{\partial Z \partial I_2} + \frac{\partial J_\mu}{\partial I_2}\frac{\partial^2 J_\nu}{\partial Z \partial I_3} - \frac{\partial J_\nu}{\partial I_2}\frac{\partial^2 J_\mu}{\partial I_3 \partial Z} - \frac{\partial J_\mu}{\partial I_3}\frac{\partial^2 J_\nu}{\partial Z \partial I_2}\right\},\end{aligned} \tag{A.19}$$

and also

$$\frac{\partial\Omega_\mu}{\partial Z} = -\frac{1}{\Delta^2}\frac{\partial\Delta}{\partial Z}\frac{\partial(J_\nu,J_\lambda)}{\partial(I_2,I_3)} + \frac{1}{\Delta}\left\{\frac{\partial J_\lambda}{\partial I_3}\frac{\partial^2 J_\nu}{\partial Z\partial I_2} + \frac{\partial J_\nu}{\partial I_2}\frac{\partial^2 J_\lambda}{\partial Z\partial I_3} - \frac{\partial J_\lambda}{\partial I_2}\frac{\partial^2 J_\nu}{\partial I_3\partial Z} - \frac{\partial J_\nu}{\partial I_3}\frac{\partial^2 J_\lambda}{\partial Z\partial I_2}\right\}, \tag{A.20}$$

and finally

$$\frac{\partial\Omega_\nu}{\partial Z} = -\frac{1}{\Delta^2}\frac{\partial\Delta}{\partial Z}\frac{\partial(J_\lambda,J_\mu)}{\partial(I_2,I_3)} + \frac{1}{\Delta}\left\{\frac{\partial J_\mu}{\partial I_3}\frac{\partial^2 J_\lambda}{\partial Z\partial I_2} + \frac{\partial J_\lambda}{\partial I_2}\frac{\partial^2 J_\mu}{\partial Z\partial I_3} - \frac{\partial J_\mu}{\partial I_2}\frac{\partial^2 J_\lambda}{\partial I_3\partial Z} - \frac{\partial J_\lambda}{\partial I_3}\frac{\partial^2 J_\mu}{\partial Z\partial I_2}\right\}. \tag{A.21}$$

The second derivatives in these expressions are evaluated by explicit differentiation of equation (A.12), etc., wrt Z. We have to perform all calculations with θ as the dummy variable, since the integrand of equation (A.8), etc., is not defined at the limits of the integration. We note that τ is now regarded as a function of θ and the integrals, and its derivative must be considered when differentiating any expression involving τ wrt an integral. We find

$$\frac{\partial^2 J_\tau}{\partial E^2} = -\frac{1}{8\pi}\oint \mathrm{d}\theta\left(\frac{1}{(\tau+\beta)\tilde{p}_\tau^2}\frac{\partial\tilde{p}_\tau}{\partial E} + \frac{1}{(\tau+\beta)^2\tilde{p}_\tau}\frac{\partial\tau}{\partial E}\right), \tag{A.22}$$

$$\frac{\partial^2 J_\tau}{\partial E\partial I_2} = \frac{1}{8\pi}\oint \frac{\mathrm{d}\theta}{(\tau+\alpha)(\tau+\beta)\tilde{p}_\tau}\left\{\frac{1}{(\tau+\alpha)}\frac{\partial\tau}{\partial E} + \frac{1}{(\tau+\beta)}\frac{\partial\tau}{\partial E} + \frac{1}{\tilde{p}_\tau}\frac{\partial\tilde{p}_\tau}{\partial E}\right\} = -\frac{1}{8\pi}\oint \frac{\mathrm{d}\theta}{(\tau+\beta)\tilde{p}_\tau}\left\{\frac{1}{(\tau+\beta)}\frac{\partial\tau}{\partial I_2} + \frac{1}{\tilde{p}_\tau}\frac{\partial\tilde{p}_\tau}{\partial I_2}\right\}, \tag{A.23}$$

$$\frac{\partial^2 J_\tau}{\partial E\partial I_3} = \frac{1}{8\pi}\oint \frac{\mathrm{d}\theta}{(\tau+\gamma)(\tau+\beta)\tilde{p}_\tau}\left\{\frac{1}{(\tau+\gamma)}\frac{\partial\tau}{\partial E} + \frac{1}{(\tau+\beta)}\frac{\partial\tau}{\partial E} + \frac{1}{\tilde{p}_\tau}\frac{\partial\tilde{p}_\tau}{\partial E}\right\} = -\frac{1}{8\pi}\oint \frac{\mathrm{d}\theta}{(\tau+\beta)\tilde{p}_\tau}\left\{\frac{1}{(\tau+\beta)}\frac{\partial\tau}{\partial I_3} + \frac{1}{\tilde{p}_\tau}\frac{\partial\tilde{p}_\tau}{\partial I_3}\right\}, \tag{A.24}$$

$$\frac{\partial^2 J_\tau}{\partial I_2^2} = \frac{1}{8\pi}\oint \frac{\mathrm{d}\theta}{(\tau+\alpha)(\tau+\beta)\tilde{p}_\tau}\left\{\frac{1}{(\tau+\alpha)}\frac{\partial\tau}{\partial I_2} + \frac{1}{(\tau+\beta)}\frac{\partial\tau}{\partial I_2} + \frac{1}{\tilde{p}_\tau}\frac{\partial\tilde{p}_\tau}{\partial I_2}\right\}, \tag{A.25}$$

$$\frac{\partial^2 J_\tau}{\partial I_2\partial I_3} = \frac{1}{8\pi}\oint \frac{\mathrm{d}\theta}{(\tau+\alpha)(\tau+\beta)\tilde{p}_\tau}\left\{\frac{1}{(\tau+\alpha)}\frac{\partial\tau}{\partial I_3} + \frac{1}{(\tau+\beta)}\frac{\partial\tau}{\partial I_3} + \frac{1}{\tilde{p}_\tau}\frac{\partial\tilde{p}_\tau}{\partial I_3}\right\} = \frac{1}{8\pi}\oint \frac{\mathrm{d}\theta}{(\tau+\gamma)(\tau+\beta)\tilde{p}_\tau}\left\{\frac{1}{(\tau+\gamma)}\frac{\partial\tau}{\partial I_2} + \frac{1}{(\tau+\beta)}\frac{\partial\tau}{\partial I_2} + \frac{1}{\tilde{p}_\tau}\frac{\partial\tilde{p}_\tau}{\partial I_2}\right\}, \tag{A.26}$$

$$\frac{\partial^2 J_\tau}{\partial I_3^2} = \frac{1}{8\pi}\oint \frac{\mathrm{d}\theta}{(\tau+\gamma)(\tau+\beta)\tilde{p}_\tau}\left\{\frac{1}{(\tau+\gamma)}\frac{\partial\tau}{\partial I_3} + \frac{1}{(\tau+\beta)}\frac{\partial\tau}{\partial I_3} + \frac{1}{\tilde{p}_\tau}\frac{\partial\tilde{p}_\tau}{\partial I_3}\right\}, \tag{A.27}$$

where we understand that everything in the integrand is evaluated at constant θ. We further observe that

$$\frac{1}{(\tau_x+\alpha)(\tau_x+\gamma)}\frac{\partial\tau_x}{\partial E} = -\frac{1}{(\tau_x+\gamma)}\frac{\partial\tau_x}{\partial I_2} = -\frac{1}{(\tau_x+\alpha)}\frac{\partial\tau_x}{\partial I_3}. \tag{A.28}$$

We now wish to evaluate $\partial\tilde{p}_\tau/\partial Z|_\theta$. We observe from equation (A.5) that

$$\tilde{p}^2\hat{\tau}^2\cos^2\theta = p^2, \tag{A.29}$$

Differentiating the above expression and rearranging, we find

$$\left.\frac{\partial \tilde{p}}{\partial Z}\right|_{\theta} = \frac{1}{2\tilde{p}\hat{\tau}^2\cos^2\theta}\left.\frac{\partial p^2}{\partial Z}\right|_{\theta} - \frac{\tilde{p}}{\hat{\tau}}\left.\frac{\partial \hat{\tau}}{\partial Z}\right|_{\theta}. \tag{A.30}$$

We note that

$$\frac{\mathrm{d}p^2}{\mathrm{d}Z} = \left.\frac{\partial p^2}{\partial Z}\right|_{\tau} + \frac{\partial \tau}{\partial Z}\left.\frac{\partial p^2}{\partial \tau}\right|_{Z}, \tag{A.31}$$

and that,

$$\left.\frac{\partial p^2}{\partial E}\right|_{\tau} = \frac{1}{2(\tau+\beta)}; \qquad \left.\frac{\partial p^2}{\partial I_2}\right|_{\tau} = -\frac{1}{2(\tau+\alpha)(\tau+\beta)}; \qquad \left.\frac{\partial p^2}{\partial I_3}\right|_{\tau} = -\frac{1}{2(\tau+\beta)(\tau+\gamma)}. \tag{A.32}$$

We also note that by differentiating equation (A.5), we find

$$\left.\frac{\partial \hat{\tau}}{\partial Z}\right|_{\theta} = \frac{1}{2\hat{\tau}\cos^2\theta}\left\{(\tau_1-\tau)\left(\left.\frac{\partial \tau}{\partial Z}\right|_{\theta} - \left.\frac{\partial \tau_0}{\partial Z}\right|_{\theta}\right) + (\tau-\tau_0)\left(\left.\frac{\partial \tau_1}{\partial Z}\right|_{\theta} - \left.\frac{\partial \tau}{\partial Z}\right|_{\theta}\right)\right\}, \tag{A.33}$$

and that $\partial\tau/\partial Z|_\theta$ can be similarly obtained as a function of $\partial\tau_0/\partial Z|_\theta$ and $\partial\tau_1/\partial Z|_\theta$ by differentiating equation (A.3). To evaluate equation (A.30), we now have to compute the derivatives of the apses $\tau_x = (\tau_0, \tau_1)$ wrt the integrals. From the definition of τ_x

$$p_\tau^2(\tau_x) = 0 = \frac{(\tau_x+\alpha)(\tau_x+\gamma)E - (\tau_x+\gamma)I_2 - (\tau_x+\alpha)I_3 + F(\tau_x)}{2(\tau_x+\alpha)(\tau_x+\beta)(\tau_x+\gamma)}. \tag{A.34}$$

Differentiating the above expression, we find

$$\frac{\partial \tau_x}{\partial E}\left((2\tau_x+\gamma+\alpha)E - I_2 - I_3 + \left.\frac{\partial F}{\partial \tau}\right|_{\tau_x}\right) = -(\tau_x+\alpha)(\tau_x+\gamma), \tag{A.35}$$

for non-trivial values of $(\alpha, \beta, \gamma, \tau_x)$. Similarly, we obtain expressions for the derivatives of the apses wrt the other integrals

$$\frac{\partial \tau_x}{\partial I_2}\left((2\tau_x+\gamma+\alpha)E - I_2 - I_3 + \left.\frac{\partial F}{\partial \tau}\right|_{\tau_x}\right) = (\tau_x+\gamma), \tag{A.36}$$

$$\frac{\partial \tau_x}{\partial I_3}\left((2\tau_x+\gamma+\alpha)E - I_2 - I_3 + \left.\frac{\partial F}{\partial \tau}\right|_{\tau_x}\right) = (\tau_x+\alpha), \tag{A.37}$$

again, for non-trivial values of $(\alpha, \beta, \gamma, \tau_x)$. We must now evaluate $\partial F/\partial\tau$. We recall that,

$$F(\tau) = (\tau+\alpha)(\tau+\gamma)G(\tau), \tag{A.38}$$

and therefore,

$$\frac{\mathrm{d}F}{\mathrm{d}\tau} = G(\tau)(\tau+\gamma) + G(\tau)(\tau+\alpha) + (\tau+\alpha)(\tau+\gamma)\frac{\mathrm{d}G}{\mathrm{d}\tau}, \tag{A.39}$$

which we can evaluate once we have chosen de Zeuuw's function $G(\tau)$. We now have expressions for all terms, and can evaluate the derivatives of the frequencies with respect to the actions.

A.4 Computing the angles

We start with the generating function for the oblate axisymmetric Stäckel potential

$$S(\tau, E, I_2, I_3) = S_\lambda + S_\phi + S_\nu, \tag{A.40}$$

where,

$$S_\tau = \int_{\tau_0}^{\tau} p_\tau \mathrm{d}\tau, \tag{A.41}$$

$$S_\phi = \int_0^{\phi} L_z \mathrm{d}\phi. \tag{A.42}$$

The generalization to the triaxial Stäckel potential is easy, since $\phi \to \nu$, and we lose equation (A.42) and gain another instance of equation (A.41) in its stead. In all cases, the angles ϑ_i are given by,

$$\vartheta_i = \frac{\partial S}{\partial J_i} = \sum_{j,k} \frac{\partial S_j}{\partial I_k} \frac{\partial I_k}{\partial J_i}, \tag{A.43}$$

where the sum in (j, k) is over each of (λ, ϕ, ν), and where I_i denote the classical integrals (E, I_2, I_3). We note that we have already calculated the $\partial I_i / \partial J_i$ in equation (A.18), etc.

In the case of S_ϕ, we can see immediately that

$$\frac{\partial S_\phi}{\partial I_i} = \begin{cases} \frac{\phi}{\sqrt{2I_2}} & \text{if } J_i = L_z, \\ 0 & \text{otherwise.} \end{cases} \tag{A.44}$$

For S_τ, we see that

$$\frac{\partial S_\tau}{\partial I_i} = \begin{cases} \frac{1}{4} \int_{\tau_0}^{\tau} \frac{\mathrm{d}\tau}{p_\tau(\tau+\beta)} & \text{if } I_i = E, \\ -\frac{1}{4} \int_{\tau_0}^{\tau} \frac{\mathrm{d}\tau}{p_\tau(\tau+\alpha)(\tau+\beta)} & \text{if } I_i = I_2, \\ -\frac{1}{4} \int_{\tau_0}^{\tau} \frac{\mathrm{d}\tau}{p_\tau(\tau+\gamma)(\tau+\beta)} & \text{if } I_i = I_3. \end{cases} \tag{A.45}$$